

FlexPod Datacenter for SAP Solution with Cisco UCS 3rd Generation Fabric and NetApp AFF A-Series

Deployment Guide for FlexPod Datacenter for SAP Solution with NetApp ONTAP 9.6 on NetApp AFF A-Series and Cisco UCS M5 Servers with 2nd Generation Intel Xeon Scalable Processors

Published: November 2020



In partnership with:  NetApp

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, go to:

<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)

© 2019 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	7
Solution Overview	8
Introduction.....	8
Audience	8
Purpose of this Document.....	8
What's New in this Release?	8
Reference.....	8
Deployment Hardware and Software	9
Architecture	9
Topology	10
Software Revisions.....	10
Configuration Guidelines	11
Physical Infrastructure	12
Network Switch Configuration.....	14
Physical Connectivity	14
FlexPod Cisco Nexus Base	14
Set Up Initial Configuration	15
FlexPod Cisco Nexus Switch Configuration.....	17
Enable Licenses	17
Create VLANs	17
Create Port Channels and assign interfaces	17
Configure Port Channel Parameters	20
Configure Virtual Port Channels	20
Uplink into Existing Network Infrastructure	21
Cisco UCS Solution for SAP HANA TDI	23
Cisco UCS Server Configuration	23
Initial Setup of Cisco UCS 6332 Fabric Interconnect	23
Cisco UCS 6332 Fabric Interconnect A.....	23
Cisco UCS 6332 Fabric Interconnect B.....	24
Log into Cisco UCS Manager.....	24
Upgrade Cisco UCS Manager Software to Version 4.0(4d).....	24
Add Block of IP Addresses for KVM Access	24
Synchronize Cisco UCS to NTP.....	25
Cisco UCS Blade Chassis Connection Options.....	25
Edit Chassis Discovery Policy.....	25
Enable Server and Uplink Ports.....	26

Acknowledge Cisco UCS Chassis and Rack-Mount Servers.....	27
Create Uplink Port Channels to Cisco Nexus Access layer Switches	27
Create New Organization	31
Create MAC Address Pools	31
Create UUID Suffix Pool.....	33
Create IQN Pools for iSCSI Boot.....	34
Create IP Pools for iSCSI Boot.....	35
Power Policy	37
Power Control Policy.....	37
Create Network Control Policy	38
Create Host Firmware Package	39
Create Local Disk Configuration Policy (Optional)	40
Create Server BIOS Policy	41
Create Serial Over LAN Policy	45
Update Default Maintenance Policy.....	46
Adapter Policy Configuration – HANA.....	47
Network Configuration	48
LAN Tab Configurations	48
Create vNIC Templates.....	54
Create vNIC/vHBA Placement Policy	64
Create LAN Connectivity Policy.....	66
Create iSCSI Boot Policy.....	73
Create Service Profile Template for SAP HANA Node	74
Configure Storage Provisioning.....	75
Configure Networking Options	75
Configure Storage Options.....	77
Configure Zoning Options	77
Configure vNIC/HBA Placement.....	77
Configure vMedia Policy.....	79
Configure Server Boot Order	80
Configure Maintenance Policy.....	85
Configure Server Assignment	85
Configure Operational Policies	86
Create Service Profiles for Single-Host System	86
Create Service Profiles for Multi-Host System.....	87
Storage Configuration	89
Complete Configuration Worksheet.....	89

Configure ONTAP Nodes	89
Configure Node 01	89
Configure Node 02	90
Set Up ONTAP Cluster	92
Set Auto-Revert on Cluster Management	97
Set Up Management Broadcast Domain	97
Create Aggregates	98
Verify Storage Failover	99
Disable Flow Control on 100GbE Ports	99
Configure AutoSupport	100
Enable Cisco Discovery Protocol	100
Create Broadcast Domains	100
Create Interface Groups	101
Create VLANs	102
Configure HTTPS Access	102
Configure SVM for the Infrastructure	104
Create SVM for the Infrastructure	106
Create Load-Sharing Mirrors	106
Create Export Policies for the Root Volumes	106
Add Infrastructure SVM Management LIF	107
Create iSCSI LIFs	107
Create Block Protocol (iSCSI) Service	107
Create FlexVol Volumes	107
Configure LUNs for iSCSI Boot	108
Configure SVM for HANA	108
Create SVM for SAP HANA	109
Create Load-Sharing Mirrors	110
Create Export Policies for the Root Volumes	110
Add HANA SVM Management Interface and Administrator	111
Create Export Policies for the HANA SVM	111
Create NFS LIF for SAP HANA Data	111
Create NFS LIF for SAP HANA Log	112
HANA Node Preparation	113
SAP HANA Node OS Preparation – RHEL for SAP HANA 7.6	113
OS Installation	113
Post Installation Tasks	122
SAP HANA Node OS Preparation – SLES for SAP 15	130

OS Installation	130
Post Installation Steps	160
System Provisioning for SAP HANA	177
Configuring SAP HANA Single-Host Systems	178
Configuration Example for a SAP HANA Single-Host System	180
Persistent Memory Configuration	183
Configure SAP HANA Multiple-Host Systems	187
Configuration Example for a SAP HANA Multiple-Host Systems	189
SAP HANA Installation	198
Important SAP Notes	198
SAP HANA IMDB Related Notes	198
Linux Related Notes	198
SAP Application Related Notes	199
Third Party Software	199
NetApp Technical reports	199
High-Availability Configuration for Multiple-host	200
SAP HANA Installation Preparations for NFSv4	200
SAP HANA Hosts	200
Storage Controllers	200
SAP HANA Data Volume Size	200
Monitor SAP HANA with AppDynamics	202
Introduction	202
SAP Landscape Monitoring	202
Trial Registration	203
Agent Installation	204
Prerequisites	204
Java Agent Installation	205
ABAP Agent Installation	205
Server Visibility Agent Installation	207
About the Authors	209
Acknowledgements	209



Executive Summary

This document describes the deployment methodology of Cisco and NetApp® FlexPod Datacenter for SAP HANA based on 2nd Generation Intel Xeon Scalable Processors supported Cisco UCS Computing System (Cisco UCS).

Cisco UCS Manager (UCSM) 4.0(4) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 and 6454), 2200/2300 series IOM, Cisco UCS B-Series Blade and Cisco UCS C-Series Rack Formfactor servers. FlexPod Datacenter with Cisco UCS unified software release 4.0(4d) and NetApp ONTAP 9.6, is a predesigned, best-practice data center architecture built on the Cisco UCS, the Cisco Nexus® 9000 family of switches, and NetApp AFF A-Series storage arrays.

Solution Overview

Introduction

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. Business agility requires application agility, so IT teams must provision applications quickly and resources must scale up (and out) as needed.

FlexPod Datacenter is a best practice data center architecture that was designed and validated by Cisco and NetApp to meet the needs of enterprise customers and service providers. FlexPod Datacenter is built on NetApp AFF enterprise storage, the Cisco UCS, and the Cisco Nexus family of switches. These components combine to create management synergy across a business's IT infrastructure. FlexPod Datacenter has been proven to be the optimal platform for a wide variety of workloads, including bare metal and virtualized systems, which enables enterprises to standardize their IT infrastructure.

Audience

The audience for this document includes, but is not limited to: field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF storage and Cisco Nexus 9000 solution.

What's New in this Release?

The following design elements distinguish this version of FlexPod from previous FlexPod models:

- Support for the Cisco UCS 4.0(4) unified software release, Cisco UCS B200-M5 servers, Cisco UCS B480-M5 servers with Cascade Lake CPUs, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Validation with cloud-scale FX Series Nexus switches
- Support for the NetApp AFF A320 storage controller
- Support for the latest release of NetApp ONTAP® 9.6 storage software
- Support for NFS v4.1
- Support for NetApp SnapCenter® 4.2
- NFS and iSCSI storage design

Reference

The design guide counterpart for this deployment guide can be accessed here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_datacenter_sap_netappaffa_design.html

Deployment Hardware and Software

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized workloads. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution.

Figure 1 shows the solution built on FlexPod components and the network connections for a configuration with the Cisco UCS 6300 series Fabric Interconnects. This design has port-channelled 40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects, port-channelled 40 Gb Ethernet connections between the C-Series rackmounts and the Cisco UCS Fabric Interconnects, and 40 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and 40/100 Gb Ethernet connection between Cisco Nexus 9000 and NetApp AFF A320 storage array. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Topology


Figure 1 Solution Topology

Cisco Unified Computing System
 UCS B-Series Blade Servers with UCS VIC 1440

Cisco UCS 6332 Fabric Interconnects

Cisco Nexus 9336C-FX2

NetApp storage controllers
 AFF-A320



The reference end-to-end 40Gb based hardware configuration includes:

- Two Cisco Nexus 9336C-FX2 switches
- Two Cisco UCS 6300 series fabric interconnects managing Cisco UCS blade/rack formfactor servers.
- One NetApp AFF A320 (HA pair) running ONTAP 9.6 with external disk shelf NS224 using NVMe SSD disks



Each Cisco UCS blade and rack server must use at least two active 40GbE ports. The blade servers must use the Port Expander Card for the Cisco VIC 1440, if this VIC is being used

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
-------	--------	-------	----------

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6300 series, UCS B480 M5, Cisco UCS C-220 M5	4.0(4d)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, Cisco UCS VIC 1440/1480 and Cisco UCS VIC 1457
Network	Cisco Nexus 9336C-FX2 NX-OS	7.0(3)I7(6)	
Storage	NetApp AFF A320	ONTAP 9.6	NetApp SnapCenter 4.2

Configuration Guidelines

This document explains how to configure 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 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Infra-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
    [-node] <nodename>                Node
    { [-vlan-name] {<netport>|<ifgrp>} VLAN Name
    | -port {<netport>|<ifgrp>}        Associated Network Port
    [-vlan-id] <integer> }            Network Switch VLAN Identifier
```

Example:

```
network port vlan create -node <node01> -vlan-name a0a-<vlan id>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. [Table 2](#) describes the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of Band Mgmt	VLAN for out-of-band management interfaces	176
Native	VLAN to which untagged frames are assigned	1

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, a cabling diagram was used.

The cabling diagram in this section contains details for the prescribed and supported configuration of the NetApp AFF A320 running NetApp ONTAP® 9.6.



For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool \(IMT\)](#).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

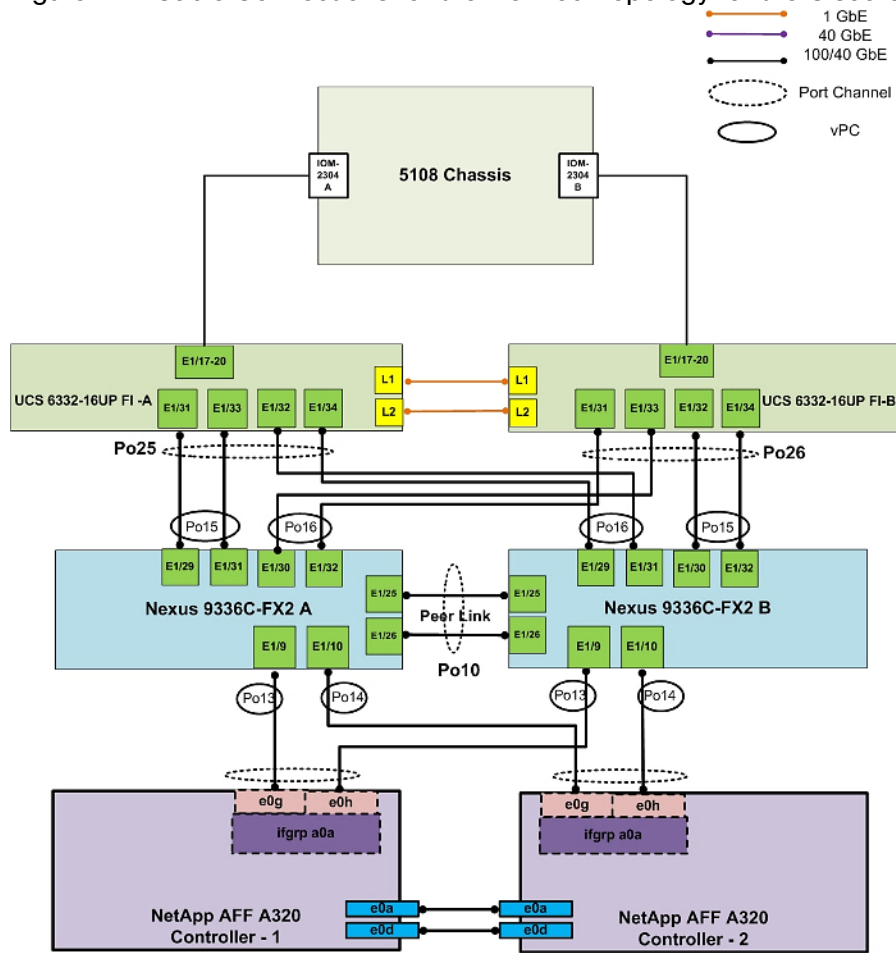


Be sure to use the cabling directions in this section as a guide.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to [NetApp Support](#).

Figure 22 illustrates the cable connections used in the validation lab for the FlexPod topology based on the Cisco UCS 6300 Series Fabric Interconnects. Also, 40Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and 40/100Gb links connect the NetApp AFF controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has two connections to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Figure 2 Cable Connections for the FlexPod Topology for the Cisco UCS 6300 Series FIs



Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as explained in the section [FlexPod Cabling](#).

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I7(6), the Cisco suggested Nexus switch release at the time of this validation.

See [Table 2](#) for the VLANs necessary for deployment as outlined in this guide.

Table 3 Configuration Variables

Variable Name	VLAN Purpose	ID Used in Validating This Document
<nexus-A-hostname>	Cisco Nexus A host name	
<nexus-A-mgmt0-ip>	Out-of-band Cisco Nexus A management IP address	
<nexus-A-mgmt0-netmask>>	Out-of-band management network netmask	
<nexus-A-mgmt0-gw>	Out-of-band management network default gateway	
<nexus-B-hostname>	Cisco Nexus Management B host name	
<nexus-B-mgmt0-ip>	Out-of-band Cisco Nexus Management B management IP address	
<nexus-B-mgmt0-netmask>>	Out-of-band management network netmask	
<nexus-B-mgmt0-gw>	Out-of-band management network default gateway	
<global-ntp-server-ip>>	NTP server IP address	
<oob-vlan-id>	Out-of-band management network VLAN ID	
<nexus-vpc-domain-id>>	Unique Cisco Nexus switch VPC domain ID	
<hana-admin-vlan-id>	Inband mgmt VLAN ID for HANA nodes for administration	76
<hana-internode-vlan-id>	HANA server-server communication network VLAN ID	220
<data-vlan-id>	HANA persistence - Data network VLAN ID	201

Variable Name	VLAN Purpose	ID Used in Validating This Document
<log-vlan-id>	HANA persistence - Log network VLAN ID	228
<backup-vlan-id>	HANA node backup VLAN	224
<hana-client-vlan-id>	Client Network for HANA VLAN ID	222
<hana-appserver-vlan-id>	Application Server Network for HANA VLAN ID	223
<hana-datasource-vlan-id>	Data source Network for HANA VLAN ID	221
<hana-replication-vlan-id>	Replication Network for HANA VLAN ID	225
<iscsi-A-vlan-id>	iSCSI-A VLAN ID	128
<iscsi-B-vlan-id>	iSCSI-B VLAN ID	129



With Cisco Nexus 9000 release 7.0(3)I7(6), autonegotiation (40G/100G) is not supported on ports 1-6 and 33-36 on the Cisco Nexus 9336C-FX2 switch. For these ports, port speed and duplex need to be hard set at both ends of the connection.

Set Up Initial Configuration

Cisco Nexus A

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

1. Configure the switch.



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

```

Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

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

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y

```

```
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus B

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

1. Configure the switch.



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

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

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

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

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

2. Review the configuration summary before enabling the configuration:

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


FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus A and Cisco Nexus B

To license the Cisco Nexus switches, follow these steps:

1. Log in as admin.
2. Run the following commands to enable the required features:

```
config t
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

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

```
vlan 76
  name HANA-Node-Mgmt
vlan 128
  name iSCSI-A
vlan 129
  name iSCSI-B
vlan 201
  name NFS_data
vlan 220
  name HANA-Internode
vlan 221
  name HANA-DataSource
vlan 222
  name HANA-Client
vlan 223
  name HANA-AppServer
vlan 224
  name HANA-Backup
vlan 225
  name HANA-System-Replication
vlan 228
  name NFS_log
exit
```

Create Port Channels and assign interfaces

Cisco Nexus A

To create the necessary port channels between devices, follow these steps:

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

```
interface Po2
```

```
description vPC peer-link

interface Eth1/25-26
channel-group 2 mode active
no shutdown

interface Po13
description PC-NetApp-A

interface Eth1/9
description AFF-A320-A:e0g
channel-group 13 mode active
no shutdown

interface Po14
description PC-NetApp-B

interface Eth1/10
description AFF-A320-B:e0g
channel-group 14 mode active
no shutdown

interface Po15
description PC1-from-FI-A

interface Eth1/29
description FI-A:1/31
channel-group 15 mode active
no shutdown

interface Eth1/31
description FI-A:1/33
channel-group 15 mode active
no shutdown

interface Po16
description PC1-from-FI-A

interface Eth1/30
description FI-B:1/31
channel-group 16 mode active
no shutdown

interface Eth1/32
description FI-B:1/33
channel-group 16 mode active
no shutdown

interface port-channel17
description PC2-from-FI-A

interface Eth1/27
description FI-A:1/27
channel-group 17 mode active
no shutdown

interface port-channel18
description PC2-from-FI-A

interface Eth1/28
description FI-B:1/27
channel-group 18 mode active
no shutdown

exit
copy run start
```

Cisco Nexus B

To create the necessary port channels between devices, follow these steps:

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

```
interface Po2
description vPC peer-link

interface Eth1/25-26
channel-group 2 mode active
no shutdown

interface Po13
description PC-NetApp-A

interface Eth1/9
description AFF-A320-A:e0h
channel-group 13 mode active
no shutdown

interface Po14
description PC-NetApp-B

interface Eth1/10
description AFF-A320-B:e0h
channel-group 14 mode active
no shutdown

interface Po15
description PC1-from-FI-A

interface Eth1/29
description FI-A:1/32
channel-group 15 mode active
no shutdown

interface Eth1/31
description FI-A:1/34
channel-group 15 mode active
no shutdown

interface Po16
description PC1-from-FI-A

interface Eth1/30
description FI-B:1/32
channel-group 16 mode active
no shutdown

interface Eth1/32
description FI-B:1/34
channel-group 16 mode active
no shutdown

interface port-channell7
description PC2-from-FI-A

interface Eth1/27
description FI-A:1/28
channel-group 17 mode active
no shutdown

interface port-channell8
description PC2-from-FI-A

interface Eth1/28
description FI-B:1/28
channel-group 18 mode active
no shutdown

exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow this step on both switches:

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

```
interface Po2
description vPC peer-link
switchport mode trunk
switchport trunk allowed vlan 76,128-129,201,220-225,228
spanning-tree port type network

interface Po13
description PC-NetApp-A
switchport mode trunk
switchport trunk allowed vlan 128-129,201,224,228
spanning-tree port type edge trunk
mtu 9216

interface Po14
description PC-NetApp-B
switchport mode trunk
switchport trunk allowed vlan 128-129,201,224,228
spanning-tree port type edge trunk
mtu 9216

interface Po15
description PC-from-FI-A
switchport mode trunk
switchport trunk allowed vlan 76,128-129,201,220-223,228
spanning-tree port type edge trunk
mtu 9216

interface Po16
description PC-from-FI-B
switchport mode trunk
switchport trunk allowed vlan 76,128-129,201,220-223,228
spanning-tree port type edge trunk
mtu 9216
exit

interface Po17
description PC2-from-FI-A
switchport mode trunk
switchport trunk allowed vlan 224-225
spanning-tree port type edge trunk
mtu 9216

interface Po18
description PC2-from-FI-B
switchport mode trunk
switchport trunk allowed vlan 224-225
spanning-tree port type edge trunk
mtu 9216
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, follow this step:

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

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po2
vpc peer-link
```

```
interface Po13
vpc 13
interface Po14
vpc 14
interface Po15
vpc 15
interface Po16
vpc 16
interface Po17
vpc 17
interface Po18
vpc 18
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, follow this step:

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

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po2
vpc peer-link
```

```
interface Po13
vpc 13
interface Po14
vpc 14
interface Po15
vpc 15
interface Po16
vpc 16
interface Po17
vpc 17
interface Po18
vpc 18
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the infrastructure. The previously described

procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

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.

Cisco UCS Server Configuration

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

See [Table 2](#) for the VLANs necessary for deployment as outlined in this guide.

Table 4 Configuration Variables

Variable Name	VLAN Purpose	ID Used in Validating This Document
<ucs-clustername>	Cisco UCS Manager cluster host name	
<ucs-cluster-ip>	Cisco UCS Manager cluster IP address	
<ucs-a-mgmt-ip>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	
<ucs-mgmt-netmask>	Out-of-band management network netmask	
<ucs_mgmt_gateway>>	Out-of-band management network default gateway	
<ucs-b-mgmt-ip>	Cisco UCS FI B out-of-band management IP address	

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, follow these steps:

1. Connect to the console port on the first Cisco UCS 6300 series Fabric Interconnect.

```

Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have chosen to setup a a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
Enter the password for "admin": <password>
Enter the same password for "admin": <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: <ucs-clustername>
Physical switch Mgmt0 IPv4 address: <ucs-a-mgmt-ip>
Physical switch Mgmt0 IPv4 netmask: <ucs-mgmt_netmask>
IPv4 address of the default gateway: <ucs-mgmt-gw>

```

```
Cluster IPv4 address: <ucsm-cluster-ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <nameserver_ip>
Configure the default domain name? y
Default domain name: <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, follow these 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: <password>
Physical switch Mgmt0 IPv4 address: <ucs-b-mgm-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 into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, follow these 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 into Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 4.0(4d)

This document assumes the use of Cisco UCS Manager Software version 4.0(4d). To upgrade the Cisco UCS Manager software and the Cisco UCS 6332 Fabric Interconnect software to version 4.0(4d), 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, follow these 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, follow these 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 <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, follow these 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.

Chassis/FEX Discovery Policy

Action :

Link Grouping Preference : None Port Channel

Backplane Speed Preference : 40G 4x10G

Enable Server and Uplink Ports

To enable server and uplink ports, follow these 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.
7. In the validation setup, Eth ports 1/19 and 1/20 on FI-A and FI-B are connected to the chassis. Right click these ports and “Configure as Server Port”.
8. Select ports that are connected to the Cisco Nexus switches, right-click them, and select “Configure as Uplink Port”.
9. In the validation setup, this is done for the eth ports 1/27-28 and 1/31-34.
10. Click Yes to confirm uplink ports and click OK.
11. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
12. Expand Ethernet Ports.
13. 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.

14. Click Yes to confirm server ports and click OK.
15. Select ports that are connected to the Cisco Nexus switches, right-click them and select Configure as Uplink Port.
16. 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, follow these 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

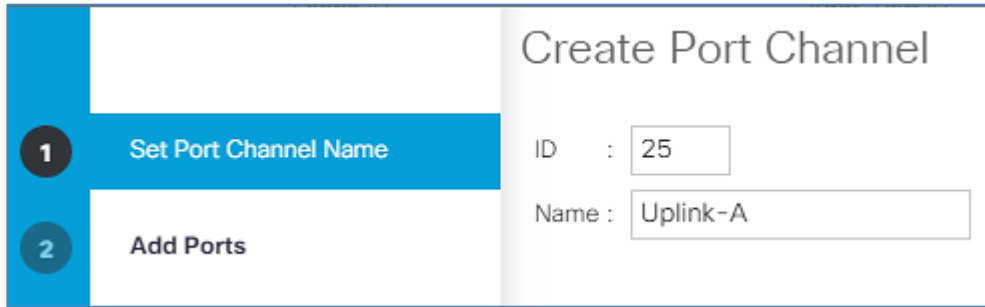
Create Uplink Port Channels to Cisco Nexus Access layer Switches

An uplink port channel 25 on Fi-A and port channel 26 on FI-B with 4 ports eth 1/31-1/34 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.

Optionally another port channel with two ports each on FI- A and FI-B can be configured for backup traffic usage, if bandwidth isolation is warranted. This 2 x 40G port channel providing 80Gbps operational speed should suffice for the same.

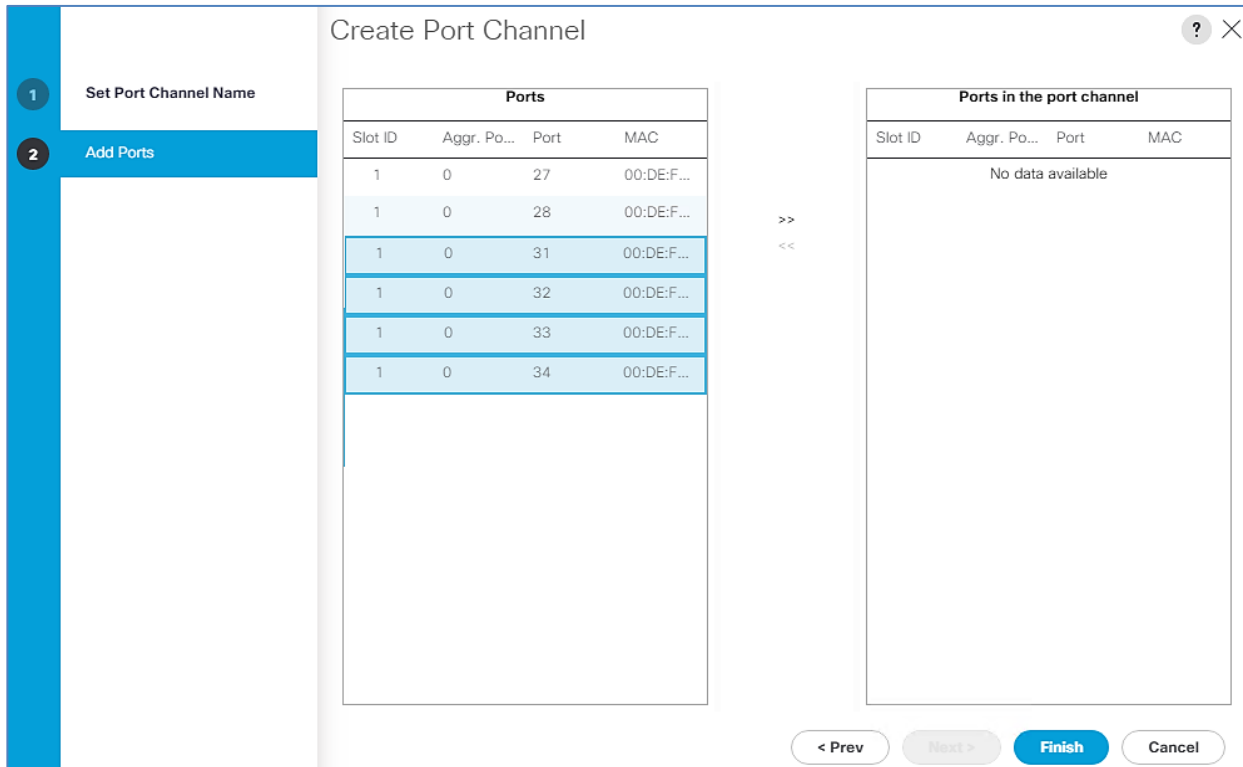
To configure the necessary port channels as planned above for the Cisco UCS environment, follow these steps:

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 25 as the unique ID of the port channel.
7. Enter Uplink-A as the name of the port channel.
8. Click Next.



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

- Slot ID 1 and port 31
- Slot ID 1 and port 32
- Slot ID 1 and port 33
- Slot ID 1 and port 34



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 26 as the unique ID of the port channel.
- d. Enter Uplink-B as the name of the port channel.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

ID :

Name :

14. Click Next.

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

- Slot ID 1 and port 31
- Slot ID 1 and port 32
- 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	27	00:DE:F...
1	0	28	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			

< Prev
Next >
Finish
Cancel

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

17. Click Finish to create the port channel.

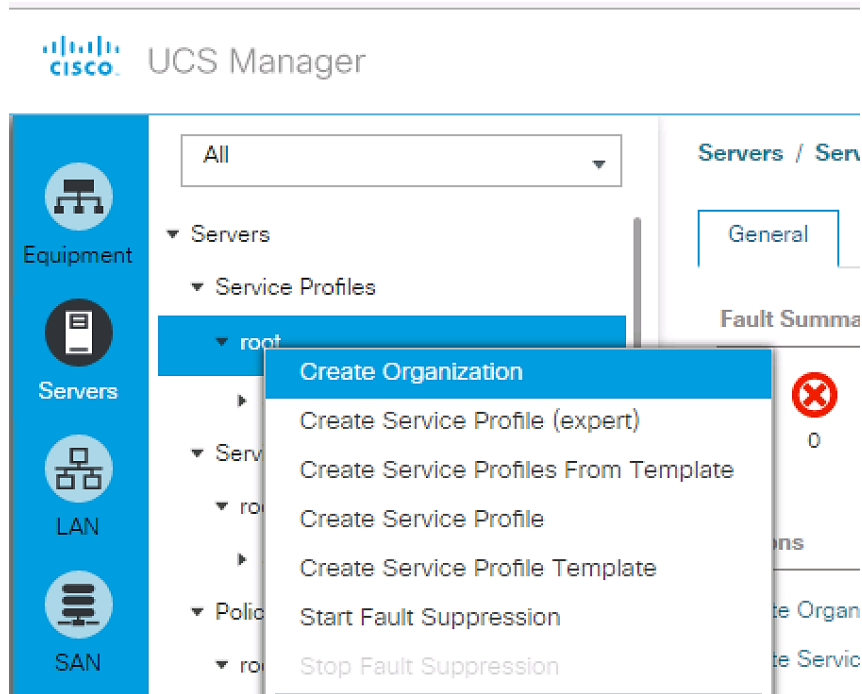
18. Click OK.
19. To create optional additional port-channel for segregated network, follow these steps:
20. In Cisco UCS Manager, click the LAN tab in the navigation pane.
21. Under LAN > LAN Cloud, expand the Fabric A tree.
22. Right-click Port Channels.
23. Select Create Port Channel
24. Enter 35 as the unique ID of the port channel.
25. Enter Uplink2-A as the name of the port channel.
26. Click Next.
27. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 27
 - Slot ID 1 and port 28
28. Click >> to add the ports to the port channel.
29. Click Finish to create the port channel.
30. Click OK.
31. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
32. Right-click Port Channels.
33. Select Create Port Channel.
34. Enter 36 as the unique ID of the port channel.
35. Enter Uplink2-B as the name of the port channel.
36. Click Next.
37. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 27
 - Slot ID 1 and port 28
38. Click >> to add the ports to the port channel.
39. Click Finish to create the port channel.
40. 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, follow these steps:

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



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

Create MAC Address Pools

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root.
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 HANA-FAB-A as the name of the MAC pool.

- Optional: Enter a description for the MAC pool.
- Choose Assignment Order Sequential.

- Click Next.
- Click Add.
- Specify a starting MAC address.
- The recommendation is to place 0A in the fourth octet of the starting MAC address to identify all the MAC addresses as Fabric Interconnect A addresses.
- Specify a size for the MAC address pool that is enough to support the available blade or server resources.

- Click OK.
- Click Finish.
- In the confirmation message, click OK.
- Right-click MAC Pools under the HANA organization.
- Select Create MAC Pool to create the MAC address pool.
- Enter HANA-FAB-B as the name of the MAC pool.
- Optional: Enter a description for the MAC pool.

Create MAC Pool

1 Define Name and Description

2 Add MAC Addresses

Name : HANA-Fab-B

Description : MAC Pool for Fabric B

Assignment Order : Default Sequential

21. Click Next.

22. Click Add.

23. Specify a starting MAC address.

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:0B:00:00 Size : 128

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

OK Cancel



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 enough 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, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.

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 enough to support the available blade or server resources.

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

Create IQN Pools for iSCSI Boot

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

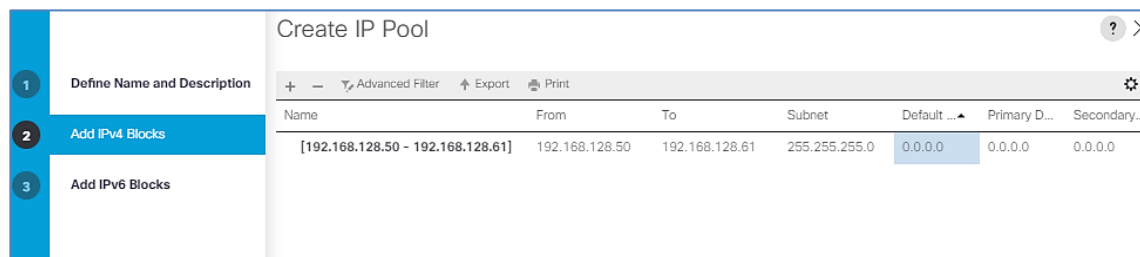
1. In Cisco UCS Manager, click SAN.
2. Expand Pools > root>Sub-Organizations>HANA.
3. Right-click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.

5. Enter IQN-HANA for the name of the IQN pool
6. Optional: Enter a description for the IQN pool
7. Enter iqn.2019-09.com.flexpod as the prefix.
8. Select Sequential for Assignment Order
9. Click Next.
10. Click Add.
11. Enter hana-node as the suffix.
12. Enter 1 in the From field.
13. Specify the size of the IQN block sufficient to support the available server resources.
14. Click OK.
15. Click Finish and OK to complete creating the IQN pool.

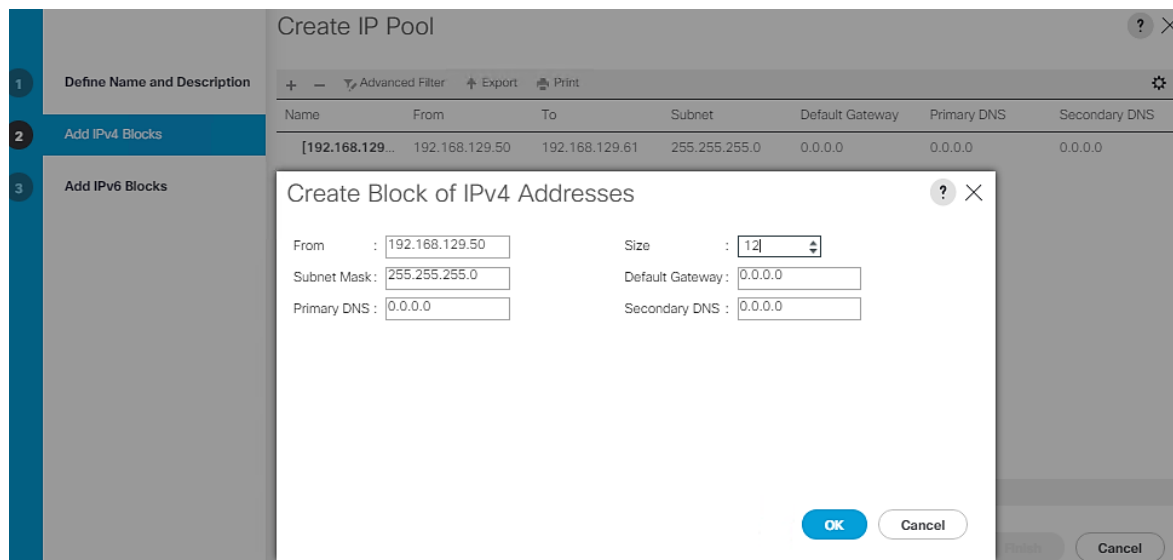
Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment within the HANA Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root > Sub-Organizations > HANA Organization.
3. Right-click IP Pools under the HANA Organization.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP addresses.



10. In the From field, enter the beginning of the range to assign as iSCSI boot IP addresses on Fabric A.
11. Set the size to enough addresses to accommodate the servers.
12. Enter the appropriate Subnet Mask.
13. Click OK.
14. Click Next.
15. Click Finish and OK to complete creating the Fabric A iSCSI IP Pool.
16. Right-click IP Pools under the FlexPod Organization.
17. Select Create IP Pool.
18. Enter iSCSI-IP-Pool-B as the name of IP pool.
19. Optional: Enter a description for the IP pool.
20. Select Sequential for the assignment order.
21. Click Next.
22. Click Add to add a block of IP addresses.



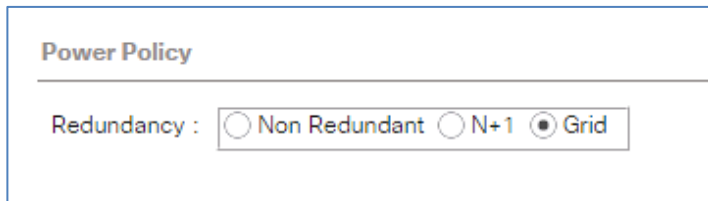
23. In the From field, enter the beginning of the range to assign as iSCSI IP addresses on Fabric B.
24. Set the size to enough addresses to accommodate the servers.
25. Enter the appropriate Subnet Mask.
26. Click OK.
27. Click Next.

28. Click Finish and OK to complete creating the Fabric B iSCSI IP Pool.

Power Policy

To run Cisco UCS with two independent power distribution units, the redundancy must be configured as Grid. Follow these 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.



Power Policy

Redundancy : Non Redundant N+1 Grid

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, follow these 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 Network Control Policy

To create a network control policy that enables CDP and LLDP on server virtual network controller (vNIC) ports, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root>Sub-Organizations>HANA.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.

Create Network Control Policy ? X

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

OK Cancel

8. Click OK to create the network control policy.

9. Click OK.

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, follow these 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 4.0(4d) packages 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
 Local Disk
 NVME Mswitch Firmware
 PSU
 Pci Switch Firmware

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, follow these 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 :

Description :

Mode :

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

FlexFlash Removable State : Yes No No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily.
Please ensure SD cards are not in use before changing the FlexFlash Removable State.

8. Click OK.

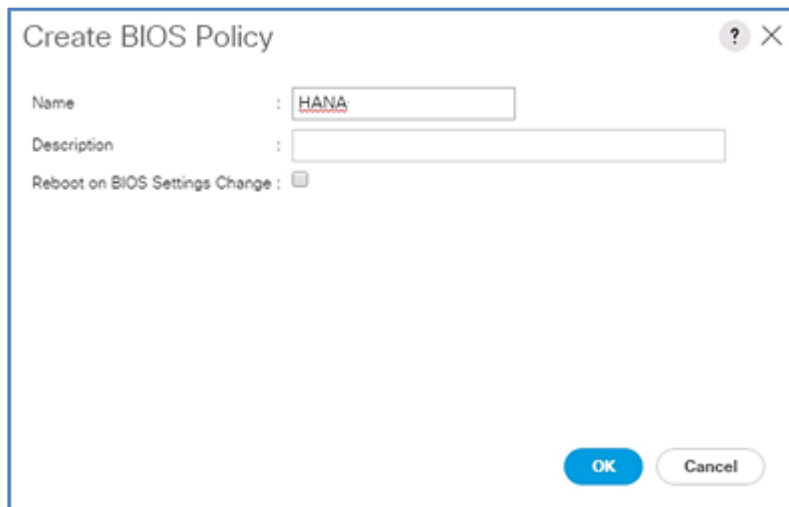
Create Server BIOS Policy



For more information, refer to the [Performance Tuning Guide for Cisco UCS M5 Servers](#) whitepaper.

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, follow these 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. Choose Create BIOS Policy.
4. Enter HANA as the BIOS policy name.
5. Click OK.



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



CPU C-States are idle power saving states. The recommendation from SAP for SAP HANA is to allow C0 and C1 states, but to disable higher C-States. This will force the CPU Core to either operate on its maximum frequency or to transition to its minimum frequency when idle.

9. Under Processor choose Custom for Power Technology and enable “C0 C1 State” for Package C State Limit option.
10. Set IO sensitive for Workload Configuration.

Servers / Policies / root / Sub-Organizations / HANA / BIOS Policies / HANA

Main | **Advanced** | Boot Options | Server Management | Events

Processor | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI

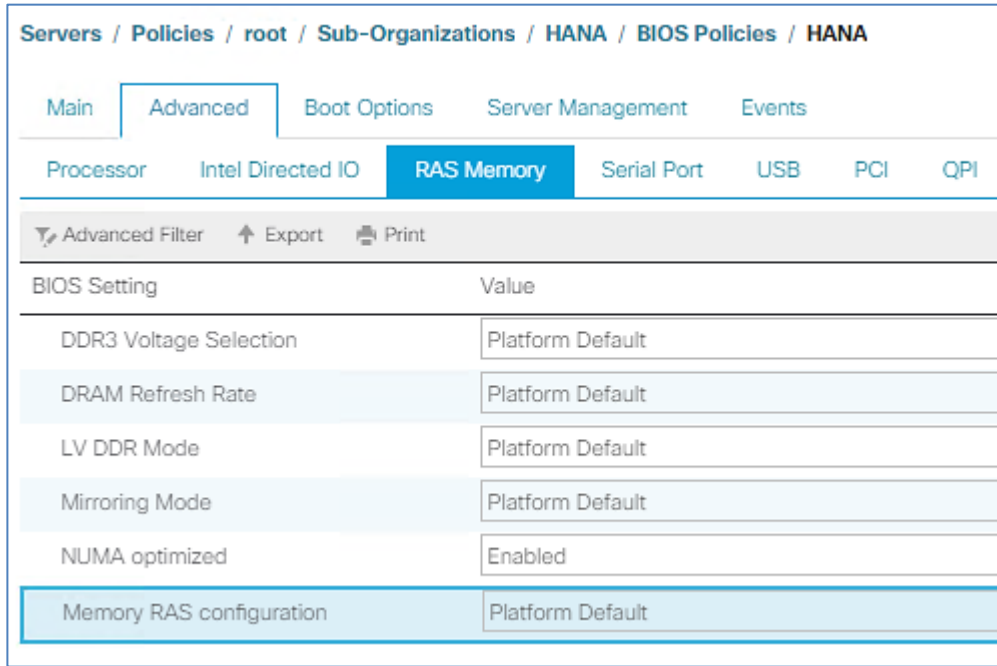
Advanced Filter | Export | Print

BIOS Setting	Value
P STATE Coordination	Platform Default
Package C State Limit	C0 C1 State
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Custom
Energy Performance	Platform Default

Workload Configuration	IO Sensitive
------------------------	--------------

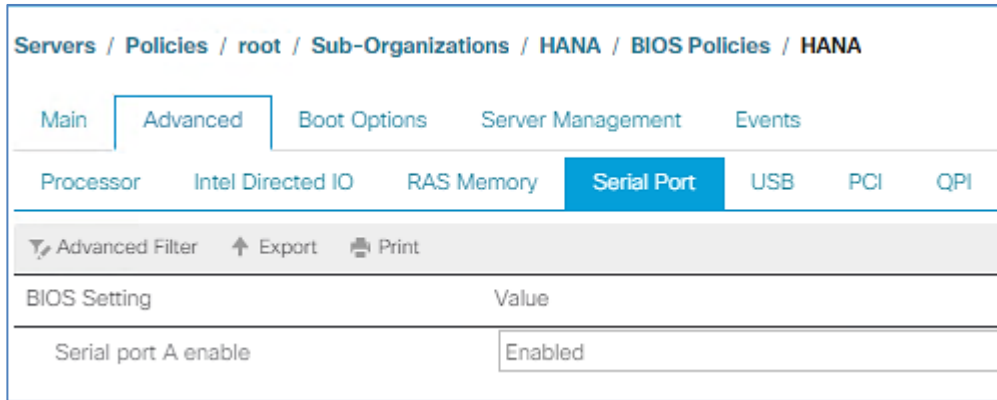
11. Click RAS Memory.

12. Choose Platform Default for Memory RAS Configuration and Enabled for NUMA optimized.



13. Click Serial Port.

14. Choose Enabled for Serial Port A enable.



15. Click Server Management.

16. 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 / Sub-Organizations / HANA / 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	Platform Default
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	Enabled
Redirection After BIOS POST	Platform Default

17. 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, follow these 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, follow these 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>

Adapter Policy Configuration – HANA

This section describes the Ethernet Adapter Policy with optimized Interrupts values. This policy must be used for the SAP HANA internal network to provide best network performance.

To create an Ethernet Adapter Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click Adapter Policies.
4. Select Create Ethernet Adapter Policy.
5. Enter HANA as the Ethernet Adapter policy name.
6. Expand Resources:
 - a. Change the Transmit Queues to 8
 - b. Change Ring size to 4096
 - c. Change the Receive Queues to 8
 - d. Change Ring size to 4096
 - e. Change the Completion Queue to 16
 - f. Change the Interrupts to 32
7. Expand Options > Change Receive Side Scaling (RSS) to Enabled
8. Change Accelerated Receive Flow Steering to Disabled
9. Click OK to create the Ethernet Adapter policy.
10. Click OK.

Servers / Policies / root / Sub-Organizations / HANA / Adapter Policies

Create Ethernet Adapter Policy

Name :

Description :

Resources

Pooled	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Transmit Queues	: <input type="text" value="8"/>	[1-1000]
Ring Size	: <input type="text" value="4096"/>	[64-4096]
Receive Queues	: <input type="text" value="8"/>	[1-1000]
Ring Size	: <input type="text" value="4096"/>	[64-4096]
Completion Queues	: <input type="text" value="16"/>	[1-2000]
Interrupts	: <input type="text" value="32"/>	[1-1024]

Options

Transmit Checksum Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Checksum Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Segmentation Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Large Receive Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Side Scaling (RSS)	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Accelerated Receive Flow Steering	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Network Virtualization using Generic Routing Encapsulation	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

OK Cancel

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 Database 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 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.

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 225 does not have to be created.

Create VLANs

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

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



In this procedure, Nine VLANs are created.

1. Select LAN > LAN Cloud.
2. Right-click VLANs.
3. Select Create VLANs.
4. Enter iSCSI_A as the name of the VLAN to be used for iSCSI -A network.
5. Keep the Common/Global option selected for the scope of the VLAN.
6. Enter <iscsi-a-vlan-id> as the ID of the HANA node admin network.
7. Keep the Sharing Type as None.
8. Click OK and then click OK again.

Create VLANs ? X

VLAN Name/Prefix :

Multicast Policy Name: [Create Multicast Policy](#)

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

VLAN IDs:

Sharing Type : None Primary Isolated Community

9. Repeat steps 1-8 to configure the rest of the VLANs as planned in the table below.

Table 5 VLANs used in this CVD

VLAN name	VLAN ID	Sample Network Address used
iSCSI_A	<iscsi-A-vlan-id>	192.168.128.x/24
iSCSI_B	<iscsi-B-valn-id>	192.168.129.x/24
Management	<hana-admin-vlan-id>	192.168.76.x/24
NFS_Data	<data-vlan-id>	192.168.201.x/24
NFS_Log	<log-vlan-id>	192.168.228.x/24
Server	<hana-internode-vlan-id>	192.168.220.x/24
DataSource	<hana-datasource-vlan-d>	192.168.221.x/24
Client	<hana-client-vlan-id>	192.168.222.x/24
AppServer	<hana-appserver-vlan-d>	192.168.223.x/24
Backup	<backup-vlan-id>	192.168.224.x/24
SysRep	<hana-replication-vlan-id>	192.168.225.x/24

Figure 3 VLANs

LAN / LAN Cloud / VLANs							
VLANs							
Advanced Filter Export Print							
Name	ID	Type	▲	Transport	Native	VLAN Sharing	
VLAN default (1)	1	Lan		Ether	Yes	None	
VLAN Management (76)	76	Lan		Ether	No	None	
VLAN IPMI (99)	99	Lan		Ether	No	None	
VLAN iSCSI_A (128)	128	Lan		Ether	No	None	
VLAN iSCSI_B (129)	129	Lan		Ether	No	None	
VLAN NFS_Data (201)	201	Lan		Ether	No	None	
VLAN Server (220)	220	Lan		Ether	No	None	
VLAN DataSource (221)	221	Lan		Ether	No	None	
VLAN Client (222)	222	Lan		Ether	No	None	
VLAN Application (223)	223	Lan		Ether	No	None	
VLAN Backup (224)	224	Lan		Ether	No	None	
VLAN SysRep (225)	225	Lan		Ether	No	None	
VLAN NFS_Log (228)	228	Lan		Ether	No	None	

Create VLAN Groups

For easier management and bandwidth segregation on the Fabric Interconnect, VLAN Groups are created within the Cisco UCS. We could bundle client zone networks, inter-node and IP storage networks in a VLAN group and create a separate VLAN group consisting of Backup and System replication networks.

VLAN group with HANA networks is assigned to 160Gbps bandwidth port channels with 4 x 40G ports such as port channels 25 and 26 on FI-A and FI-B respectively. VLAN group of backup and replication networks is assigned to 80Gbps bandwidth capable port channels with 2 x 40G ports, in other words, port channels 35 and 36 on FI-A and FI-B respectively.

To configure the necessary VLAN Groups for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLAN Groups.
4. Select Create VLAN Groups.
5. Enter HANA-nws as the name of the VLAN Group used for this bundle of networks.

6. Select all the created VLANs except for backup and system replication

Create VLAN Group

Name:

VLANs

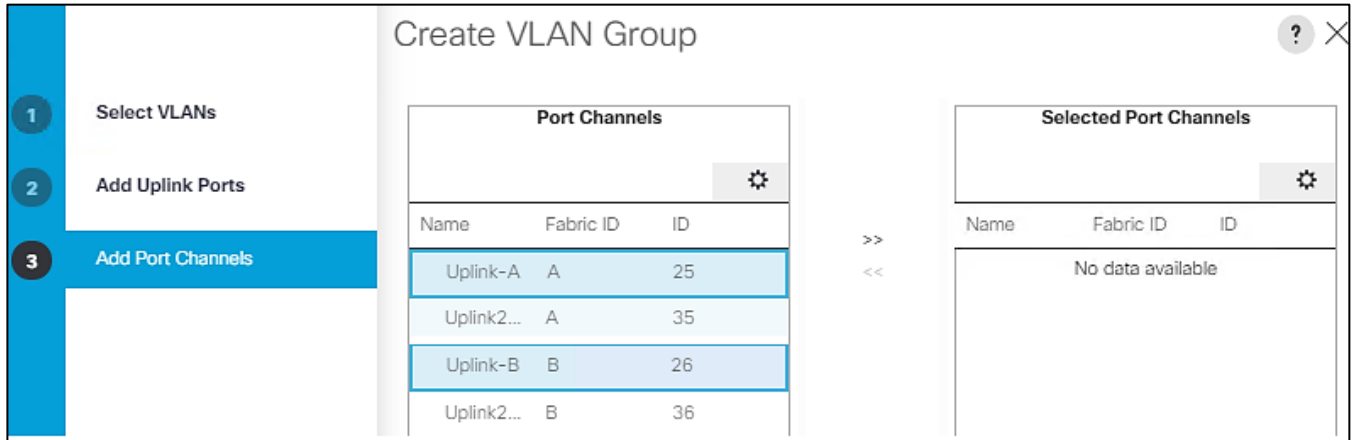
Advanced Filter Export Print No Native VLAN

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	Application	<input type="radio"/>
<input type="checkbox"/>	Backup	<input type="radio"/>
<input checked="" type="checkbox"/>	Client	<input type="radio"/>
<input checked="" type="checkbox"/>	DataSource	<input type="radio"/>
<input checked="" type="checkbox"/>	ISCSI_A	<input type="radio"/>
<input checked="" type="checkbox"/>	ISCSI_B	<input type="radio"/>
<input checked="" type="checkbox"/>	Management	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS_Data	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS_Log	<input type="radio"/>
<input checked="" type="checkbox"/>	Server	<input type="radio"/>
<input type="checkbox"/>	SysRep	<input type="radio"/>

Create VLAN

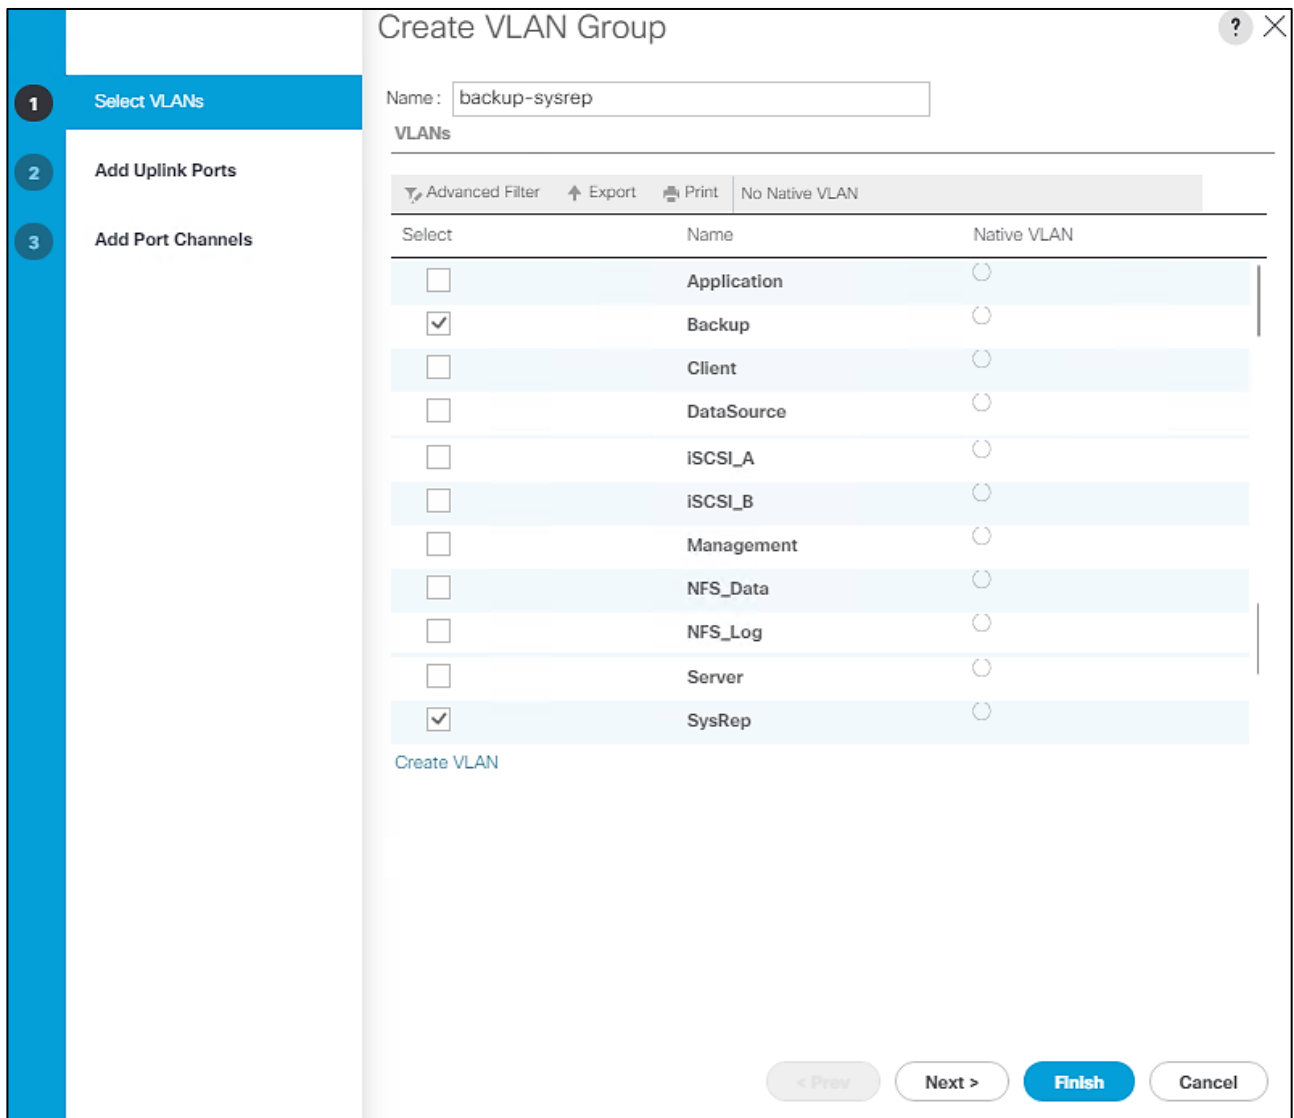
< Prev Next > Finish Cancel

7. Click Next.
8. Click Next on Add Uplink Ports, since you will use port-channel.
9. Choose port-channels created for uplink network. Click >>.



10. Click Finish.

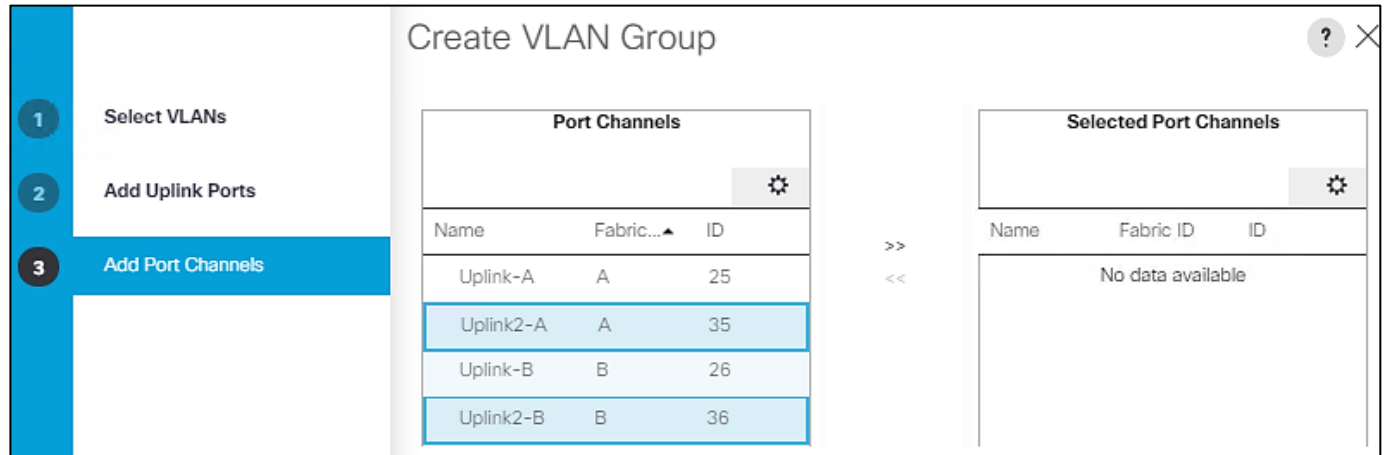
11. Similarly, create *bkp-sysrep* vlan-group including the Backup and System Replication networks.



12. Click Next.

13. Click Next on Add Uplink Ports, since you will use port-channel.

14. Choose port-channels created for uplink network. Click >>.



15. Click Finish.

Create vNIC Templates

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.

Create vNIC Template for iSCSI via Fabric A

To create iSCSI virtual network interface card (vNIC) templates for the Cisco UCS environment within the FlexPod Organization, follow these steps:

1. Select LAN.
2. Expand Policies > root > Sub-Organizations > FlexPod Organization.
3. Right-click vNIC Templates under the HANA Organization.
4. Select Create vNIC Template.
5. Enter iSCSI-A as the vNIC template name.
6. Select Fabric A. Do not select the Enable Failover checkbox.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter checkbox is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only iSCSI_A.
11. Select iSCSI_A as the native VLAN.

12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select HANA-Fab-A.
15. From the Network Control Policy list, select Enable-CDP-LLDP.



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

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

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 | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN	VLAN ID
<input checked="" type="checkbox"/>	ISCSI_A	<input checked="" type="radio"/>	128
<input type="checkbox"/>	ISCSI_B	<input type="radio"/>	129
<input type="checkbox"/>	Management	<input type="radio"/>	76
<input type="checkbox"/>	NFS_Data	<input type="radio"/>	201
<input type="checkbox"/>	NFS_hanashared	<input type="radio"/>	130

Create VLAN

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 iSCSI with Fabric B

To create a vNIC template for iSCSI with fabric B, follow these steps:

1. Right-click vNIC Templates.

2. Select Create vNIC Template.
3. Enter iSCSI-Template-B as the vNIC template name.
4. Select Fabric B. Do not select the Enable Failover checkbox.
5. Leave Redundancy Type set at No Redundancy.
6. Under Target, make sure that only the Adapter checkbox is selected.
7. Select Updating Template for Template Type.
8. Under VLANs, select only iSCSI-B.
9. Select iSCSI-B as the native VLAN.
10. Leave vNIC Name set for the CDN Source.
11. Under MTU, enter 9000.
12. From the MAC Pool list, select HANA-Fab-B.
13. From the Network Control Policy list, select Enable-CDP-LLDP.
14. Click OK to complete creating the vNIC template.
15. Click OK.

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 | VLAN Groups

Advanced Filter | Export | Print ⚙

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	ISCSI_A	<input type="radio"/>	128
<input checked="" type="checkbox"/>	ISCSI_B	<input checked="" type="radio"/>	129
<input type="checkbox"/>	Management	<input type="radio"/>	76
<input type="checkbox"/>	NFS_Data	<input type="radio"/>	201
<input type="checkbox"/>	NFS_hanashared	<input type="radio"/>	130

Create VLAN

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 :



Internal Network requires > 9.5 Gbps for SAP HANA inter-node communication; choose Platinum QoS Policy created for HANA-Internal vNIC template.

Create a vNIC Template for Inter-node Network

To create a vNIC template for Inter-node network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter Server as the vNIC template name.
6. Keep Fabric B selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for Server.
11. Set Server as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-B.
14. For Network Control Policy, choose default from drop-down list
15. Click OK to create the vNIC template.

Create a vNIC Template for NFS-Data Network

To create a vNIC template for NFS-Data network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter NFS-Data as the vNIC template name.
6. Keep Fabric A selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.

9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for NFS_Data.
11. Set NFS_Data as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-A.
14. For Network Control Policy, choose Enable-CDP-LLDP from drop-down list
15. Click OK to create the vNIC template.

Create a vNIC Template for NFS-Log Network

To create a vNIC template for an NFS-Log network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter NFS-Log as the vNIC template name.
6. Keep Fabric A selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for NFS_Log.
11. Set NFS_Log as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-B.
14. For Network Control Policy, choose Enable-CDP-LLDP from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for Application Network

To create a vNIC template for an application network, follow these steps:

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

2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter Application as the vNIC template name.
6. Keep Fabric B selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for Application.
11. Set Application as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-B.
14. For Network Control Policy, choose default from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for DataSource Network

To create vNIC template for a DataSource network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter DataSource as the vNIC template name.
6. Keep Fabric A selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for DataSource.
11. Set DataSource as the native VLAN.

12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-A.
14. For Network Control Policy, choose default from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for Backup Network

To create vNIC template for a backup network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates
4. Choose Create vNIC Template.
5. Enter Backup as the vNIC template name.
6. Keep Fabric B selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for Backup.
11. Set Backup as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-B.
14. For Network Control Policy, choose Enable-CDP-LLDP from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for System Replication Network

To create a vNIC template for a system replication network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates
4. Choose Create vNIC Template.

5. Enter SysRep as the vNIC template name.
6. Keep Fabric A selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for SysRep.
11. Set SysRep as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose HANA-Fab-A.
14. For Network Control Policy, choose default from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for Management Network

To create a vNIC template for a management network, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter Mgmt as the vNIC template name.
6. Keep Fabric A selected.
7. Check the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is unchecked.
9. Choose Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for Management.
11. Set Management as the native VLAN.
12. For MTU, enter 1500.
13. In the MAC Pool list, choose HANA-Fab-A.
14. For Network Control Policy, choose default from drop-down list.

15. Click OK to create the vNIC template.

LAN / Policies / root / Sub-Organizations / HANA / vNIC Templates		
vNIC Templates		
+ - ⌵ Advanced Filter ⬆ Export 🖨 Print		
Name	VLAN	Native VLAN
▼ vNIC Template Access		
Network Access	Access	<input checked="" type="radio"/>
▼ vNIC Template Application		
Network Application	Application	<input checked="" type="radio"/>
▼ vNIC Template Backup		
Network Backup	Backup	<input checked="" type="radio"/>
▼ vNIC Template DataSource		
Network DataSource	DataSource	<input checked="" type="radio"/>
▼ vNIC Template iSCSI-A		
Network iSCSI_A	iSCSI_A	<input checked="" type="radio"/>
▼ vNIC Template iSCSI-B		
Network iSCSI_B	iSCSI_B	<input checked="" type="radio"/>
▼ vNIC Template Mgmt		
Network Management	Management	<input checked="" type="radio"/>
▼ vNIC Template NFS-Data		
Network NFS_Data	NFS_Data	<input checked="" type="radio"/>
▼ vNIC Template NFS-Log		
Network NFS_Log	NFS_Log	<input checked="" type="radio"/>
▼ vNIC Template Server		
Network Server	Server	<input checked="" type="radio"/>
▼ vNIC Template SysRep		
Network SysRep	SysRep	<input checked="" type="radio"/>

Create vNIC/vHBA Placement Policy

Cisco UCS assigns virtual network interface connections (vCons) to the PCIe adapter cards in the server. Each vCon is a virtual representation of a physical adapter that can be assigned vNICs and vHBAs.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme. This can be one of the following:

- Round Robin– In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon3 to Adapter1, then assigns vCon2 and vCon4 to Adapter2.
- In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1, vCon2 and vCon4 to Adapter2, and vCon3 to Adapter3.
- This is the default scheme.
- Linear Ordered– In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon2 to Adapter1, then assigns vCon3 and vCon4 to Adapter2.
- In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1 and vCon2 to Adapter2, then assigns vCon3 and vCon4 to Adapter3.

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



Cisco UCS B480 M5 server used in the validation setup is configured with VIC1440 w/ port expander and VIC1480. These appear as Adpater1 and Adapter 3.

Equipment / Chassis / Chassis 2 / Servers / Server 1 / Adapters						
Adapters						
Advanced Filter Export Print						
Name	Vendor	PID	Serial	Overall Status	Operability	
Adapter 1	Cisco Systems Inc	UCSB-MLOM-40G-04	FCH22437XHR	↑ Operable	↑ Operable	
Adapter 3	Cisco Systems Inc	UCSB-VIC-M84-4P	FCH22337UJY	↑ Operable	↑ Operable	

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 and select Liner Ordered for Virtual Slot Mapping Scheme.

Servers / Policies / root / Sub-Organizations / HANA / vNIC/vHBA Placement Policies / HANA

General Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : HANA
Owner : Local
Virtual Slot Mapping Scheme : Round Robin Linear Ordered

Virtual Host Interfaces

Advanced Filter Export Print

Virtual Slot	Selection Preference	Transport
1	Assigned Only	ethernet,fc
2	Assigned Only	ethernet,fc
3	Assigned Only	ethernet,fc
4	Assigned Only	ethernet,fc

Create LAN Connectivity Policy

With LAN connectivity policy we are pre-defining the vNICs that a system needs to have in order to cater to the specific networks the use-case demands.

For both the single host system and multiple-host use-cases, apart from the admin/management the node may need backup and application server network connections, at a minimum. It could have other networks depending on the usage.

However, for the multiple-host scenario, inter-node communication is mandatory and that is what distinguishes it from a single host system's policy.

In the following steps we create separate LAN connectivity policies for the two standard use-cases. This further simplifies the service profile template creation by having the pre-determined vNICs already part of the network configuration.

Single-Host Systems Use-Case

1. To configure the necessary LAN Connectivity Policy within the HANA Organization, follow these steps:
2. In Cisco UCS Manager, click LAN.
3. Expand LAN > Policies > root > Sub-Organizations > HANA Organization.
4. Right-click LAN Connectivity Policies under the HANA Organization.
5. Select Create LAN Connectivity Policy.
6. Enter iSCSI-Boot-SciUp as the name of the policy.

7. Click Add to add a vNIC.
8. In the Create vNIC dialog box, enter HANA-node-admin as the name of the vNIC.
9. Select the Use vNIC Template checkbox.
10. In the vNIC Template list, select Mgmt.
11. In the Adapter Policy list, select Linux.
12. Click OK to add this vNIC to the policy.
13. Click Add to add another vNIC to the policy.
14. In the Create vNIC box, enter HANA-data as the name of the vNIC.
15. Select the Use vNIC Template checkbox.
16. In the vNIC Template list, select NFS-Data.
17. In the Adapter Policy list, select HANA.
18. Click OK to add the vNIC to the policy.
19. Click Add to add another vNIC to the policy.
20. In the Create vNIC dialog box, enter HANA-log as the name of the vNIC.
21. Select the Use vNIC Template checkbox.
22. In the vNIC Template list, select NFS-Log.
23. In the Adapter Policy list, select HANA.
24. Click OK to add this vNIC to the policy.
25. Click Add to add a vNIC.
26. In the Create vNIC dialog box, enter HANA-AppServer as the name of the vNIC.
27. Select the Use vNIC Template checkbox.
28. In the vNIC Template list, select Application.
29. In the Adapter Policy list, select Linux.
30. Click OK to add this vNIC to the policy.
31. Click Add to add a vNIC.
32. In the Create vNIC dialog box, enter HANA-node-backup as the name of the vNIC.

33. Select the Use vNIC Template checkbox.
34. In the vNIC Template list, select Backup.
35. In the Adapter Policy list, select Linux.
36. Click OK to add this vNIC to the policy.
37. Click Add to add a vNIC.
38. In the Create vNIC dialog box, enter iSCSI-A as the name of the vNIC.
39. Select the Use vNIC Template checkbox.
40. In the vNIC Template list, select iSCSI-A.
41. In the Adapter Policy list, select Linux.
42. Click OK to add this vNIC to the policy.
43. Click Add to add a vNIC to the policy.
44. In the Create vNIC dialog box, enter iSCSI-B as the name of the vNIC.
45. Select the Use vNIC Template checkbox.
46. In the vNIC Template list, select iSCSI-B.
47. In the Adapter Policy list, select Linux.
48. Click OK to add this vNIC to the policy.
49. Expand Add iSCSI vNICs.
50. Select Add in the Add iSCSI vNICs section.
51. Set the name to iSCSI-Boot-A.
52. Select iSCSI-A as the Overlay vNIC.
53. Set the iSCSI Adapter Policy to default.
54. Leave the VLAN set to iSCSI_A (native).
55. Leave the MAC Address set to None.
56. Click OK.
57. Select Add in the Add iSCSI vNICs section.
58. Set the name to iSCSI-Boot-B.

59. Select iSCSI-B as the Overlay vNIC.
60. Set the iSCSI Adapter Policy to default.
61. Leave the VLAN set to iSCSI_B (native).
62. Leave the MAC Address set to None
63. Click OK.

LAN / Policies / root / Sub-Organizations / HANA / LAN Connectivity Policies / **iSCSI-Boot-ScUp**

General Events

Actions

Delete

Show Policy Usage

Use Global

Name : **iSCSI-Boot-ScUp**

Description:

Owner : **Local**

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

Name	MAC Address	Native VLAN
▶ vNIC HANA-AppServer	Derived	
▶ vNIC HANA-data	Derived	
▶ vNIC HANA-log	Derived	
▶ vNIC HANA-node-admin	Derived	
▶ vNIC HANA-node-backup	Derived	
▶ vNIC iSCSI-A	Derived	
▶ vNIC iSCSI-B	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-A	iSCSI-A	default	Derived
iSCSI vNIC iSCSI-Boot-B	iSCSI-B	default	Derived

Multiple-Host Systems Use-Case

To configure the necessary LAN Connectivity Policy within the HANA Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > HANA Organization.
3. Right-click LAN Connectivity Policies under the HANA Organization.
4. Select Create LAN Connectivity Policy.
5. Enter iSCSI-Boot-ScOut as the name of the policy.

6. Click Add to add a vNIC.
7. In the Create vNIC dialog box, enter HANA-node-admin as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Mgmt.
10. In the Adapter Policy list, select Linux.
11. Click OK to add this vNIC to the policy.
12. Click Add to add another vNIC to the policy.
13. In the Create vNIC box, enter HANA-data as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select NFS-Data.
16. In the Adapter Policy list, select HANA.
17. Click OK to add the vNIC to the policy.
18. Click Add to add a vNIC.
19. In the Create vNIC dialog box, enter HANA-log as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select NFS-Log.
22. In the Adapter Policy list, select HANA.
23. Click OK to add this vNIC to the policy.
24. Click Add to add a vNIC.
25. In the Create vNIC dialog box, enter HANA-AppServer as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select Application.
28. In the Adapter Policy list, select Linux.
29. Click OK to add this vNIC to the policy.
30. Click Add to add a vNIC.
31. In the Create vNIC dialog box, enter HANA-node-backup as the name of the vNIC.

32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select Backup.
34. In the Adapter Policy list, select Linux.
35. Click OK to add this vNIC to the policy.
36. Click Add to add a vNIC.
37. In the Create vNIC dialog box, enter HANA-internode as the name of the vNIC.
38. Select the Use vNIC Template checkbox.
39. In the vNIC Template list, select Server.
40. In the Adapter Policy list, select HANA.
41. Click OK to add this vNIC to the policy
42. Click Add to add a vNIC.
43. In the Create vNIC dialog box, enter iSCSI-A as the name of the vNIC.
44. Select the Use vNIC Template checkbox.
45. In the vNIC Template list, select iSCSI-A.
46. In the Adapter Policy list, select Linux.
47. Click OK to add this vNIC to the policy.
48. Click Add to add a vNIC to the policy.
49. In the Create vNIC dialog box, enter iSCSI-B as the name of the vNIC.
50. Select the Use vNIC Template checkbox.
51. In the vNIC Template list, select iSCSI-B.
52. In the Adapter Policy list, select Linux.
53. Click OK to add this vNIC to the policy.
54. Expand Add iSCSI vNICs.
55. Select Add in the Add iSCSI vNICs section.
56. Set the name to iSCSI-Boot-A.
57. Select iSCSI-A as the Overlay vNIC.

58. Set the iSCSI Adapter Policy to default.
59. Leave the VLAN set to iSCSI_A (native).
60. Leave the MAC Address set to None.
61. Click OK.
62. Select Add in the Add iSCSI vNICs section.
63. Set the name to iSCSI-Boot-B.
64. Select iSCSI-B as the Overlay vNIC.
65. Set the iSCSI Adapter Policy to default.
66. Leave the VLAN set to iSCSI_B (native).
67. Leave the MAC Address set to None
68. Click OK

LAN / Policies / root / Sub-Organizations / HANA / LAN Connectivity Policies / iSCSI-Boot-ScOut

General | Events

Actions

- Delete
- Show Policy Usage
- Use Global

Name : **iSCSI-Boot-ScOut**

Description:

Owner : **Local**

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

Name	MAC Address	Native VLAN
▶ vNIC HANA-data	Derived	
▶ vNIC HANA-internode	Derived	
▶ vNIC HANA-log	Derived	
▶ vNIC HANA-node-admin	Derived	
▶ vNIC HANA-node-backup	Derived	
▶ vNIC iSCSI-A	Derived	
▶ vNIC iSCSI-B	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-A	iSCSI-A		Derived
iSCSI vNIC iSCSI-Boot-B	iSCSI-B		Derived

Create iSCSI Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi_lif_1a and iscsi_lif_1b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif_2a and iscsi_lif_2b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).



One boot policy is configured in this procedure. The policy configures the primary target to be iscsi_lif_1a.

To create a boot policy for the Cisco UCS environment within the HANA Organization, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > HANA Organization.
3. Right-click Boot Policies under the HANA Organization.
4. Select Create Boot Policy.
5. Enter iSCSI-HANA as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Leave Enforce on vNIC/vHBA/iSCSI Name checkbox selected.
9. Select the Legacy Mode.
10. Expand the Local Devices drop-down menu and select Add CD/DVD.
11. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
12. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
13. Click OK.
14. Select Add iSCSI Boot.
15. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.
16. Click OK.

Servers / Policies / root / Sub-Organizations / HANA / Boot Policies / Boot Policy iSCSI-HANA

General Events

Actions	Properties
Delete	Name : iSCSI-HANA
Show Policy Usage	Description : <input type="text"/>
Use Global	Owner : Local
	Reboot on Boot Order Change : <input type="checkbox"/>
	Enforce vNIC/vHBA/iSCSI Name : <input checked="" type="checkbox"/>
	Boot Mode : <input checked="" type="radio"/> Legacy <input type="radio"/> Uefi

Warning

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

Local Devices		Boot Order	
<input type="checkbox"/> Local Devices		+ - Advanced Filter Export Print	
<input type="checkbox"/> CIMC Mounted vMedia		Name	vNIC/vHBA/iSCSI...
<input type="checkbox"/> vNICs		CD/DVD	1
<input type="checkbox"/> vHBAs		▼ iSCSI	2
<input type="checkbox"/> iSCSI vNICs		iSCSI	iSCSI-Boot-A Primary
<input type="checkbox"/> EFI Shell		iSCSI	iSCSI-Boot-B Secondary
		<input type="button" value="Move Up"/> <input type="button" value="Move Down"/> <input type="button" value="Delete"/>	
		<input type="button" value="Set Uefi Boot Parameters"/>	

17. Click OK then click OK again to create the policy.

Create Service Profile Template for SAP HANA Node

This section details the service profile template creation procedure. The steps to create service profile template to instantiate HANA nodes for single-host or multiple-host use-case depends on the LAN connectivity policy you select and corresponding placement of vNICs per vCONS.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > HANA Organization.
3. Right-click the HANA Organization.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.

5. Enter iSCSI-HANA-node as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID Assignment, select UUID_Pool.
8. Click Next.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy tab select the default Local Storage Policy.
2. Click Next.

Configure Networking Options

To configure the network options for HANA node intended to be ScaleUp system, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select iSCSI-Boot-SclUp from the LAN Connectivity Policy drop-down list.
4. Select IQN-HANA in Initiator Name Assignment.

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: ▼
[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy: ▼ [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: ▼

Initiator Name :
[Create IQN Suffix Pool](#)
 The IQN will be assigned from the selected pool.
 The available/total IQNs are displayed after the pool name.

< Prev Next > **Finish** Cancel

5. Click Next.

OR

To configure the network options for HANA node intended to be part of multiple-host cluster, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select iSCSI-Boot-ScIOut from the LAN Connectivity Policy drop-down list.
4. Select IQN-HANA in Initiator Name Assignment.

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy: Create LAN Connectivity Policy

Initiator Name

Initiator Name Assignment:

Initiator Name:

[Create IQN Suffix Pool](#)

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

< Prev Next > **Finish** Cancel

5. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

1. Select No vHBAs for the “How would you like to configure SAN connectivity?” field.
2. Click Next.

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

To configure the vNIC/HBA placement in case of Single system, follow these steps:

1. In the “Select Placement” list, Select HANA.
2. Assign the vNICs to vCon1 and vCon3 as below:



Even though seven networks were defined, they are optional and if they are not needed in your deployment, the addition of a vNIC template for that network may be omitted.

Create Service Profile Template

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [Create Placement Policy](#)

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Sele...	Tran...
▼ vCon 1		Assi...	ethe...
vNIC HANA-AppServer	4		
vNIC HANA-node-admin	2		
vNIC HANA-node-backup	3		
vNIC ISCSI-A	1		
vCon 2		Assi...	ethe...

↑ Move Up ↓ Move Down

Create Service Profile Template

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [Create Placement Policy](#)

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Sele...	Tran...
vCon 2		Assi...	ethe...
▼ vCon 3		Assi...	ethe...
vNIC HANA-data	2		
vNIC HANA-log	3		
vNIC ISCSI-B	1		
vCon 4		Assi...	ethe...

↑ Move Up ↓ Move Down

3. Click Next.

To configure the vNIC/HBA placement in case of Multiple-host system, follow these steps:

1. In the “Select Placement” list, Select HANA.
2. Assign the vNICs to vCon1 and vCon3 as below:



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 that network may be omitted.

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [Create Placement Policy](#)

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Selection...	Transport
▼ vCon 1		Assigned...	ethernet.fc
vNIC HANA-AppServer	4		
vNIC HANA-internode	1		
vNIC HANA-node-admin	2		
vNIC iSCSI-A	3		
vCon 2		Assigned...	ethernet.fc

↑ Move Up ↓ Move Down

OK **Cancel**

Specify how vNICs and vHBAs are placed on physical network adapters

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Or...	Selec...	Trans...
▼ vCon 3		Assig...	ether...
vNIC HANA-data	1		
vNIC HANA-log	2		
vNIC HANA-node-backup	4		
vNIC iSCSI-B	3		
vCon 4		Assig...	ether...

↑ Move Up ↓ Move Down

OK **Cancel**

3. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

1. Do not select a vMedia Policy.
2. Click Next.

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. Select Boot-iSCSI-A for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Boot-iSCSI-A** [Create Boot Policy](#)

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

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

Boot Order

Name	vNIC/vHBA/iSCSI vNIC	Type	LI	VI	SI	BI	DI
Remote CD/DVD	1						
▼ iSCSI	2						
iSCSI	iSCSI-Boot-A	Primary					
iSCSI	iSCSI-Boot-B	Secondary					
CIMC Mounted CD/DVD	3						

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

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

2. In the Boot order, select iSCSI-Boot-A.
3. Click Set iSCSI Boot Parameters.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
5. Leave the "Initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI-IP-Pool-A as the "Initiator IP address Policy."

Set iSCSI Boot Parameters

Name: **ISCSI-Boot-A**

Authentication Profile: [Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment:

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy:

IPv4 Address : **0.0.0.0**
 Subnet Mask : **255.255.255.0**
 Default Gateway: **0.0.0.0**
 Primary DNS : **0.0.0.0**
 Secondary DNS : **0.0.0.0**

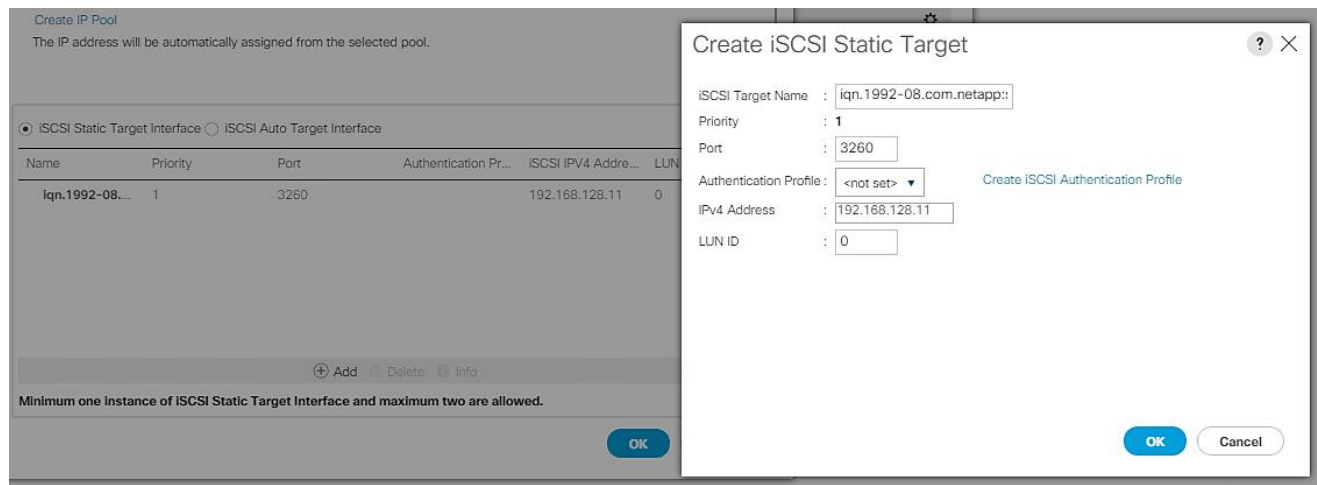
[Create IP Pool](#)

The IP address will be automatically assigned from the selected pool.

7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, log into the storage cluster management interface and run the “iscsi show” command”.

```
sap-hana::> iscsi show
      Target                Target                Status
Vserver  Name                    Alias                 Admin
-----
infra_svm iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5
                                     infra_svm             up
sap-hana::> █
```

10. Enter the IP address of iscsi_lif_1a for the IPv4 Address field.



11. Click OK to add the iSCSI static target.
12. Click Add.
13. Enter the iSCSI Target Name.
14. Enter the IP address of iscsi_lif_2a for the IPv4 Address field.
15. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters ? >

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-A(12/12) ▾

IPv4 Address : **0.0.0.0**
 Subnet Mask : **255.255.255.0**
 Default Gateway : **0.0.0.0**
 Primary DNS : **0.0.0.0**
 Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.1992-08...	1	3260		192.168.128.11	0
iqn.1992-08...	2	3260		192.168.128.12	0

⊕ Add 🗑 Delete ℹ Info

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

OK
Cancel

16. Click OK to complete setting the iSCSI Boot Parameters.

17. In the Boot order, select iSCSI-Boot-B.

The screenshot displays the Cisco UCS management interface. On the left, the 'Server Boot Order' table is visible, showing the following entries:

Name	vNIC/vHBA/iSCSI
Remote CD/DVD	1
ISCSI	2
ISCSI	ISCSI-Boot-A
ISCSI	ISCSI-Boot-B
CIMC Mounted CD/DVD	3

The 'Set iSCSI Boot Parameters' dialog box is open on the right, showing the following configuration:

- Name: **ISCSI-Boot-B**
- Authentication Profile: <not set> (with a 'Create iSCSI Authentication Profile' link)
- Initiator Name: (empty field)
- Initiator Name Assignment: <not set>
- Initiator IP Address Policy: **ISCSI-IP-Pool-B(12/12)**
- IPv4 Address: **0.0.0.0**
- Subnet Mask: **255.255.255.0**
- Default Gateway: **0.0.0.0**
- Primary DNS: **0.0.0.0**
- Secondary DNS: **0.0.0.0**
- Target Interface: ISCSI Static Target Interface ISCSI Auto Target Interface

18. Click Set iSCSI Boot Parameters.
19. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
20. Leave the "Initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
21. Set iSCSI-IP-Pool-B as the "Initiator IP address Policy".
22. Select the iSCSI Static Target Interface option.
23. Click Add.
24. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, login into storage cluster management interface and run "iscsi show" command".
25. Enter the IP address of iscsi_lif_1b for the IPv4 Address field.
26. Click OK to add the iSCSI static target.
27. Click Add.
28. Enter the iSCSI Target Name.
29. Enter the IP address of iscsi_lif_2b for the IPv4 Address field.
30. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters ? >

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI-IP-Pool-B(12/12) ▼

IPv4 Address : **0.0.0.0**
 Subnet Mask : **255.255.255.0**
 Default Gateway : **0.0.0.0**
 Primary DNS : **0.0.0.0**
 Secondary DNS : **0.0.0.0**

[Create IP Pool](#)
 The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPv4 Addre...	LUN Id
iqn.1992-08....	1	3260		192.168.129.11	0
iqn.1992-08....	2	3260		192.168.129.12	0

⊕ Add 🗑 Delete ℹ Info

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

OK
Cancel

31. Click OK to complete setting the iSCSI Boot Parameters. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.
2. Click Next.

Configure Server Assignment

To configure the server assignment, follow these steps:

1. Select Assign Later for the Pool Assignment

- Under Firmware Management select HANA-FW for Host Firmware Package from the dropdown list. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

- In the BIOS Policy list, select HANA.
- Expand Management IP Address: On Outband IPv4 tab, select Outband-Mgmt policy from the drop-down list for Management IP Address Policy.
- Expand Power Control Policy Configuration and select HANA in the Power Control Policy list.
- Choose default from drop-down list for the rest.

Create Service Profile Template

Optionally specify information that affects how the system operates.

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

External IPMI/Redfish Management Configuration
 Management IP Address
 Monitoring Configuration (Thresholds)

Power Control Policy Configuration
 Power control policy determines power allocation for a server in a given power group.
 Power Control Policy: [Create Power Control Policy](#)

Scrub Policy
 KVM Management Policy
 Graphics Card Policy
 Persistent Memory Policy

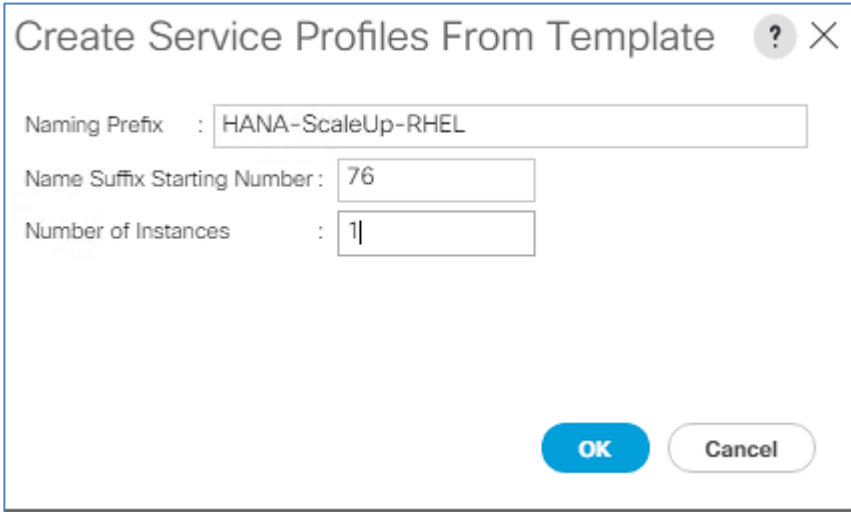
- Click Finish to create the service profile template.
- Click OK in the confirmation message.

Create Service Profiles for Single-Host System

To create service profiles from the service profile template created using the LAN connectivity policy iSCSI-Boot-ScUp, intending to create a single host system based on RHEL7.6, follow these steps:

- Connect to Cisco UCS Manager and click Servers.

2. Select Service Profile Templates > root > Sub-Organizations > HANA Organization > Service Template iSCSI-HANA-node
3. Right-click iSCSI-HANA-node and select Create Service Profiles from Template.
4. Enter HANA-ScaleUp-RHEL as the service profile prefix.
5. Enter 76 as “Name Suffix Starting Number.”
6. Enter 1 as the “Number of Instances.”



Create Service Profiles From Template ? X

Naming Prefix : HANA-ScaleUp-RHEL

Name Suffix Starting Number : 76

Number of Instances : 1

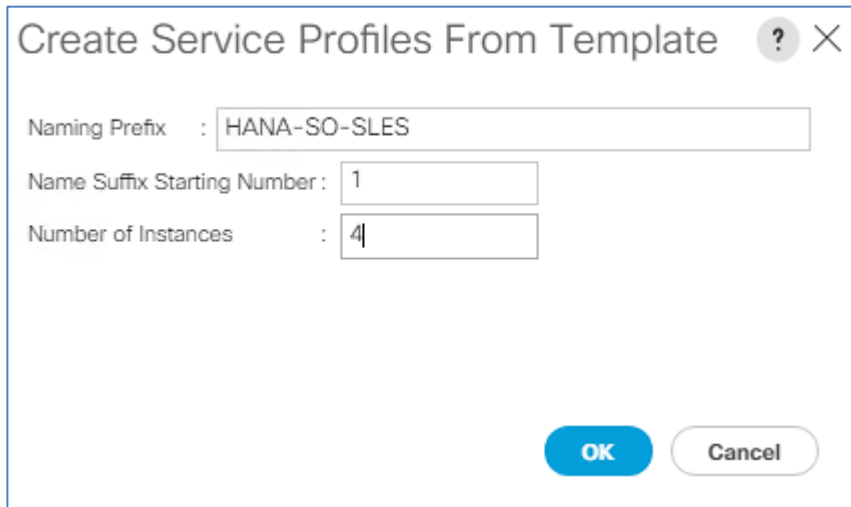
OK Cancel

7. Click OK to create the service profile.
8. Click OK in the confirmation message.

Create Service Profiles for Multi-Host System

To create service profiles from the service profile template created using the LAN connectivity policy iSCSI-Boot-ScOut, intending to create a 4-node scale-out cluster based on SLES 15, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organizations > HANA Organization > Service Template iSCSI-HANA-node
3. Right-click iSCSI-HANA-node and select Create Service Profiles from Template.
4. Enter HANA-SO-SLES as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 4 as the “Number of Instances.”



The image shows a dialog box titled "Create Service Profiles From Template" with a help icon and a close icon. It contains three input fields: "Naming Prefix" with the value "HANA-SO-SLES", "Name Suffix Starting Number" with the value "1", and "Number of Instances" with the value "4". At the bottom right, there are two buttons: "OK" (a blue button) and "Cancel" (a white button with a grey border).

7. Click OK to create the service profiles.
8. Click OK in the confirmation message.

Storage Configuration

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet in the [ONTAP 9.6 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 9.6 Software Setup Guide](#) to learn about configuring ONTAP software. Table 6 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 6 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 9.6 URL	<url-boot-software>

Configure Node 01

To configure node 01, follow these 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 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 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, follow these 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 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 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 username, indicating no username.
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 ONTAP Cluster

Table 7 lists all the parameters required to set up the ONTAP cluster.

Table 7 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 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>

Cluster Detail	Cluster Detail Value
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>

To set up an ONTAP cluster, follow these steps:

1. 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 9.6 software boots on the node for the first time.
2. 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>

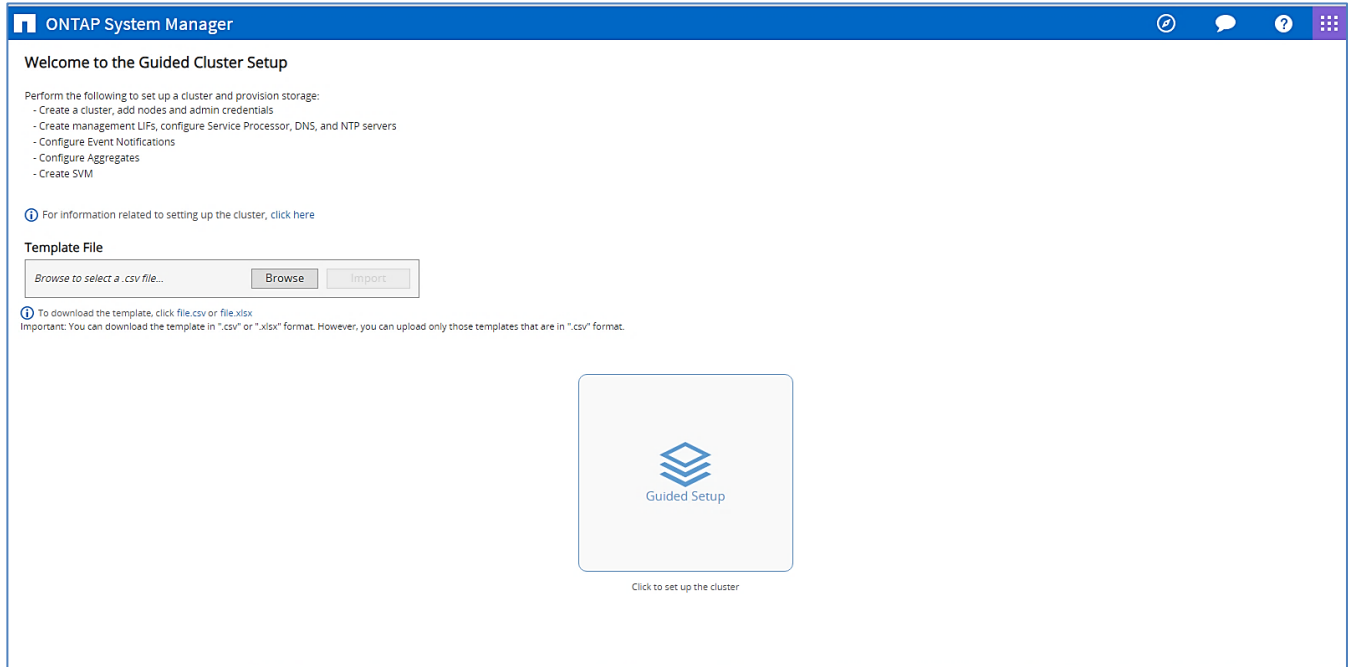
Otherwise press Enter to complete cluster setup using the command line interface:

```

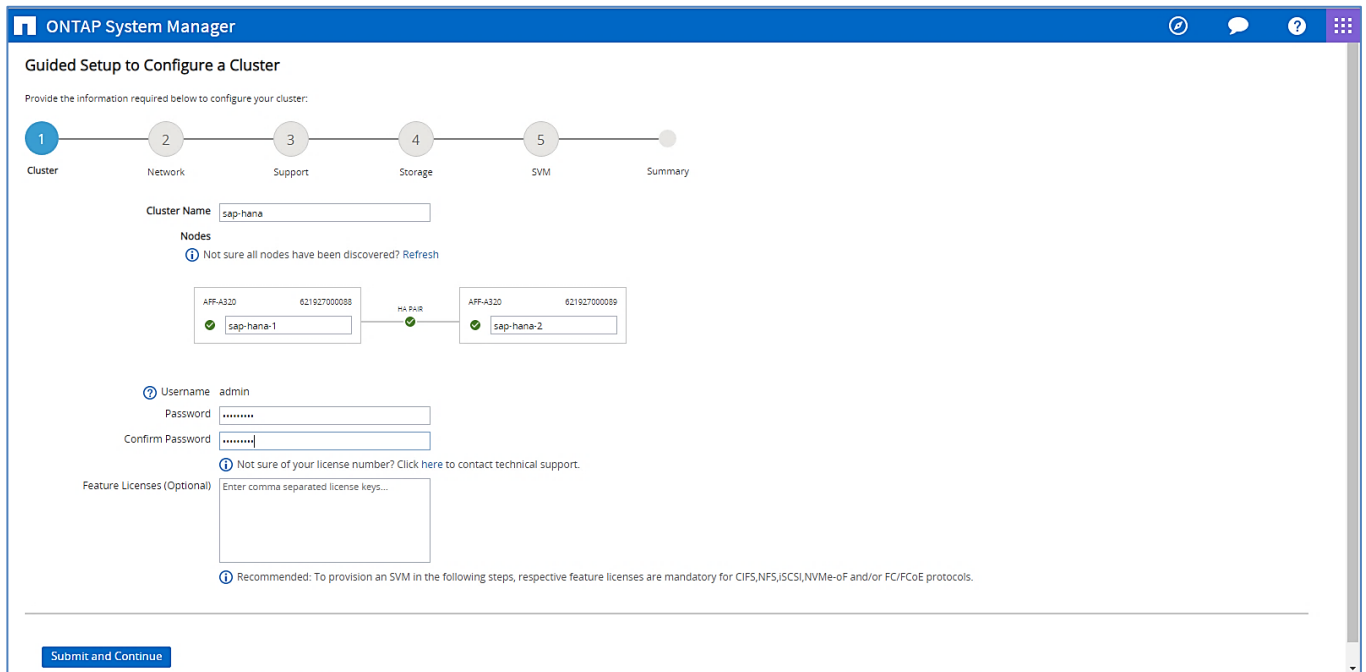


Cluster setup can also be done using the CLI guided setup. This document describes the cluster setup using NetApp System Manager.

- To complete the cluster setup, access <https://<node01-mgmt-ip>> with your web browser Start the “Guided Cluster Setup.”



- Provide the cluster name, the password of the user admin, and add all relevant licenses. Click “Submit and Continue”



- Provide the IP address, netmask, and gateway for the Cluster management and choose the desired port, for example select e0M for the first controller. Retain the settings for the Node management by checking the box.
- Provide the IP addresses for the Service Processor Management. Check the boxes “Override the Default Values” and “Retain Networks and Gateway Configuration of the Cluster Management”. Finally enter the DNS and NTP details.



The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet.

ONTAP System Manager

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster **2** Network Support Storage SVM Summary

Network (Management)

IP Addresses (IPv4) required Enter 1 Cluster Management, 0 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range You must enter the default network details manually.

Cluster Management IP Address: 192.168.76.31 Netmask: 255.255.255.0 Gateway (Optional): 192.168.76.1 Port: sap-hana-1a

Node Management Retain Netmask and Gateway configuration of the Cluster Management.

sap-hana-1 IP Address: 192.168.76.39 Netmask: Gateway: Port:

sap-hana-2 IP Address: 192.168.76.30 Netmask: Gateway: Port:

Service Processor Management Default values have been detected for the Service Processor.

Override the default values (Gateway is mandatory).

Retain Netmask and Gateway configuration of the Cluster Management.

sap-hana-1 IP Address: 192.168.76.34

sap-hana-2 IP Address: 192.168.76.35

DNS Details

DNS Domain Names:

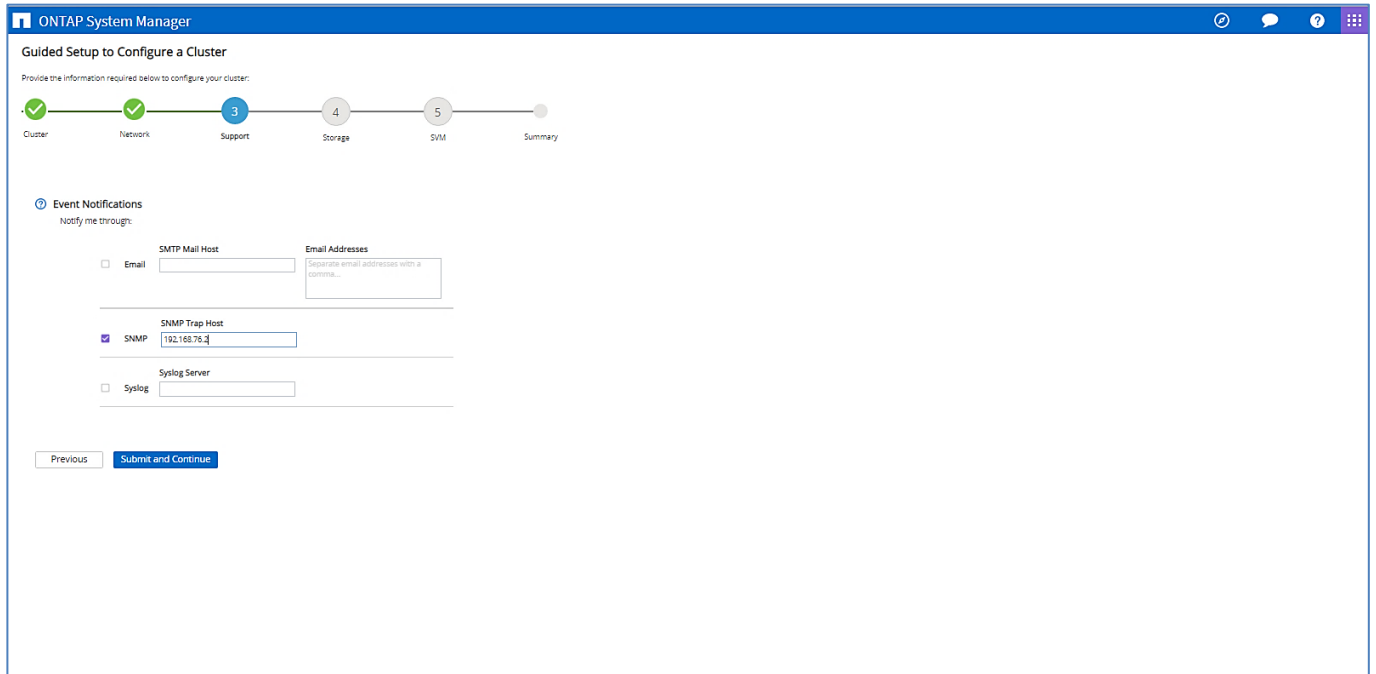
DNS Server IP Address:

NTP Details

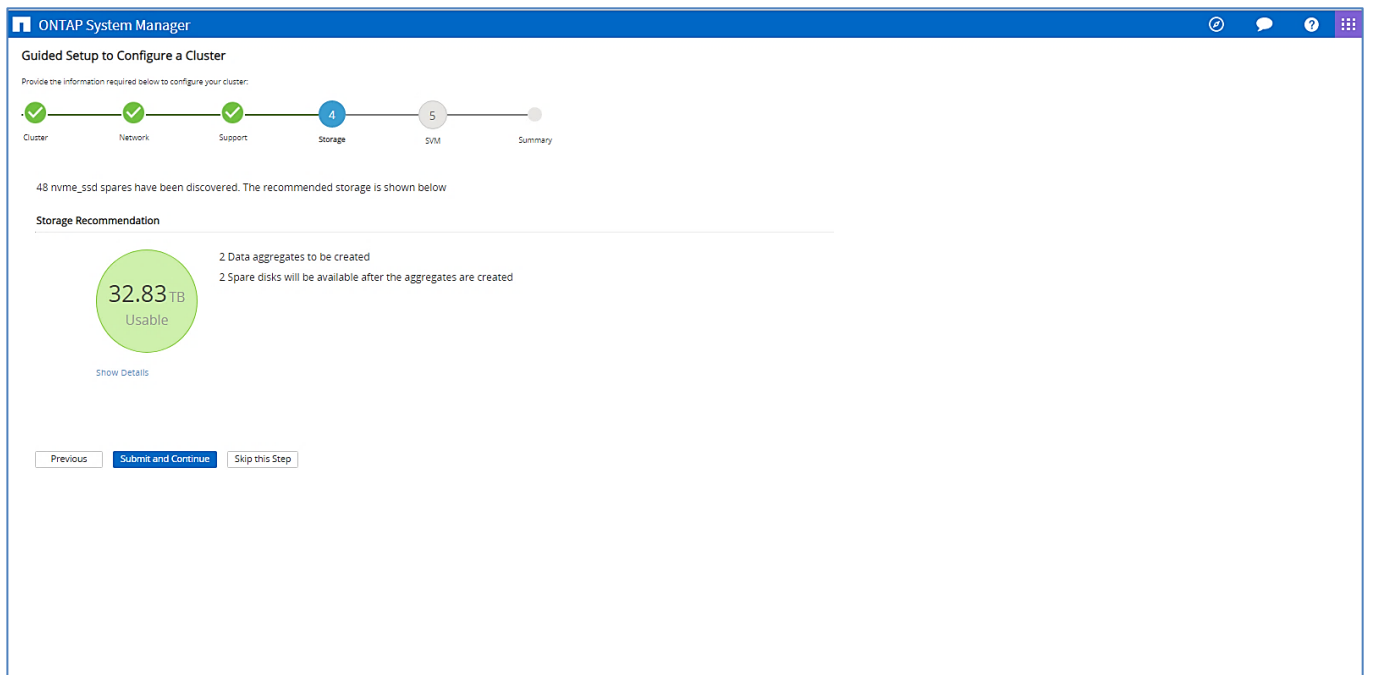
Primary NTP Server:

Alternative NTP Server (Optional):

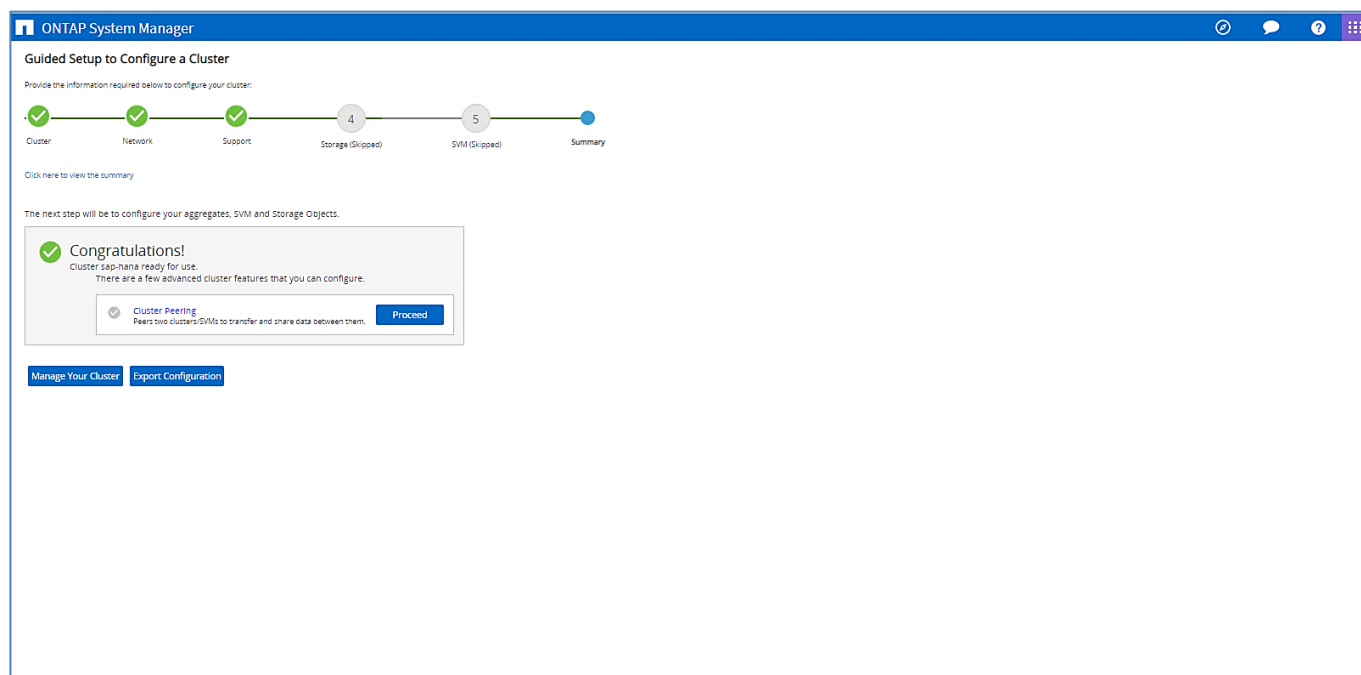
- Enter the information for any needed event notifications.



8. Skip the fourth step (Storage) without creating any aggregate. You will perform this step manually later.



9. The Cluster has been successfully configured.



To log into the cluster, follow these 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 vserver) 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 executed for each storage node. Storage nodes are named after the cluster name with an appended number.

```
broadcast-domain remove-ports -broadcast-domain Default -ports <clustername>-1:e0g,<clustername>-1:e0h,
<clustername>-2:e0g,<clustername>-2:e0h
broadcast-domain show
```

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 a high availability pair.



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, follow these steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_1 -node <clustername>-1 -diskcount 11
aggr create -aggregate aggr1_2 -node <clustername>-2 -diskcount 11
aggr create -aggregate aggr2_1 -node <clustername>-1 -diskcount 11
aggr create -aggregate aggr2_2 -node <clustername>-2 -diskcount 11
```



Use all disks except for two spares to create the aggregates. In this example 11 disks per aggregate were used.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until all are online.

- (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:

- Verify the status of storage failover.

```
storage failover show
```



Both `<clustername>-1` and `<clustername>-2` must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

- Enable failover on one of the two nodes.

```
storage failover modify -node <clustername>-1 -enabled true
```



Enabling failover on one node enables it for both nodes.

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



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

- Continue with step 6 if high availability is configured.
- 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
```

- 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 100GbE Ports

NetApp recommends disabling flow control on all the 100GbE ports that are connected to external devices. To disable flow control, follow these steps:

1. Run the following commands to configure node 1:

```
network port modify -node <clustername>-01 -port e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 2:

```
network port modify -node <clustername>-2 -port e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show -fields flowcontrol-admin
```

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 4 shows the physical network connection and the virtual LANs (VLANs) used for this setup.

Figure 4 Physical Network Connection and VLANs

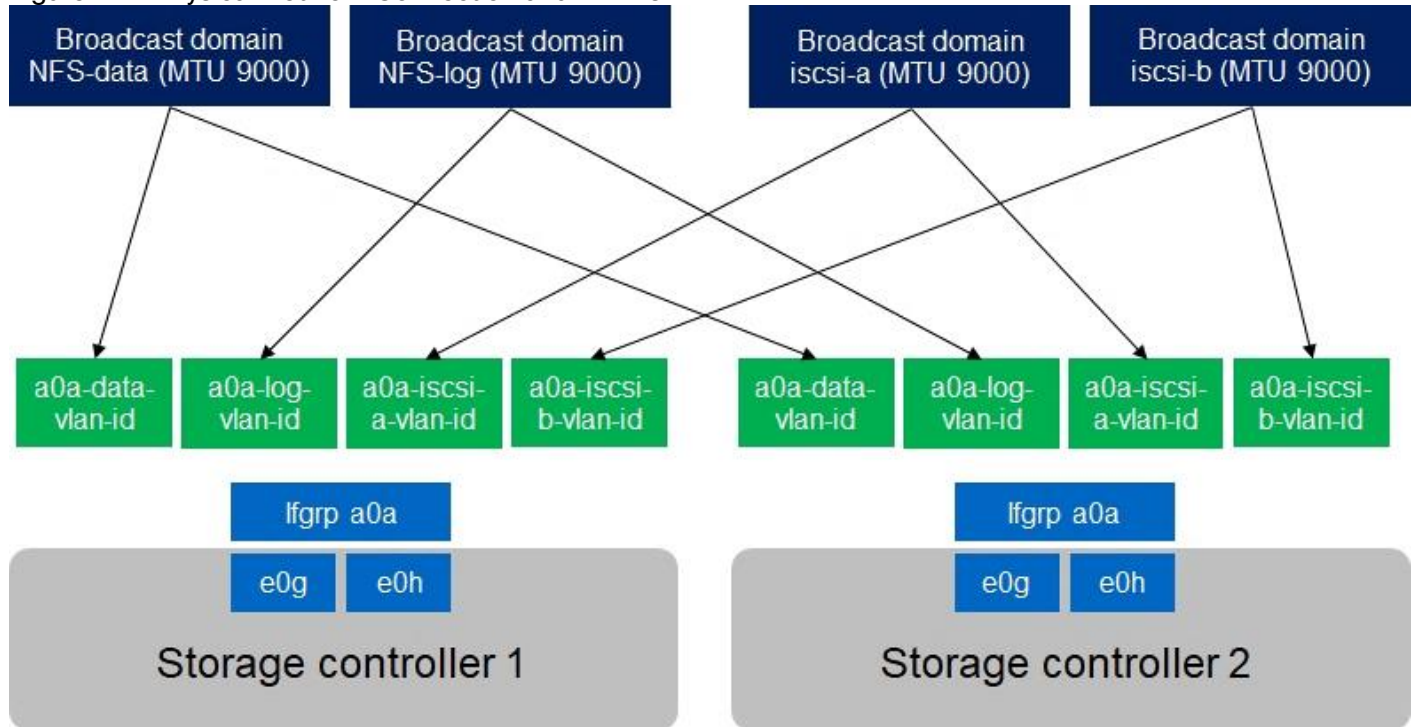


Table 8 Cluster Networking Requirements

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
iSCSI a VLAN ID	<iscsi-a-vlan-id>	128
iSCSI b VLAN ID	<iscsi-b-vlan-id>	129
Storage backend VLAN ID	<stbackend-vlan-id>	224

All broadcast domains 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 iSCSI-a -mtu 9000
broadcast-domain create -broadcast-domain iSCSI-b -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>-1 -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <clustername>-1 -ifgrp a0a -port e0g
ifgrp add-port -node <clustername>-1 -ifgrp a0a -port e0h

ifgrp create -node <clustername>-2 -ifgrp a0a -distr-func port -mode multimode_lacp
    
```

```
ifgrp add-port -node <clustername>-2 -ifgrp a0a -port e0g
ifgrp add-port -node <clustername>-2 -ifgrp a0a -port e0h

ifgrp show
```

Create VLANs

To create VLANs, follow these steps:

1. Set the MTU size of the interface groups.

```
network port modify -node <clustername>-1 -port a0a -mtu 9000
network port modify -node <clustername>-2 -port a0a -mtu 9000
```

2. Create HANA data VLAN ports and add them to the NFS-Data broadcast domain.

```
network port vlan create -node <clustername>-1 -vlan-name a0a-<data-vlan-id>
network port vlan create -node <clustername>-2 -vlan-name a0a-<data-vlan-id>

broadcast-domain add-ports -broadcast-domain NFS-Data -ports <clustername>-1: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 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 storage-backend -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 9 .

Table 9 ONTAP Software Parameter to Enable HTTPS

Cluster Detail	Cluster Detail Value
----------------	----------------------

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, follow these steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example the <serial-number>) by running the following command:

```
security certificate show
```

3. 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.

4. 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>
```

5. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the security certificate show command.
6. 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>
```

7. 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.

8. 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
```

Configure SVM for the Infrastructure

Table 10 and Figure 5 describe the infrastructure SVM together with all required storage objects (volumes, export-policies, and LIFs).

Figure 5 Overview of Infrastructure SVM Components

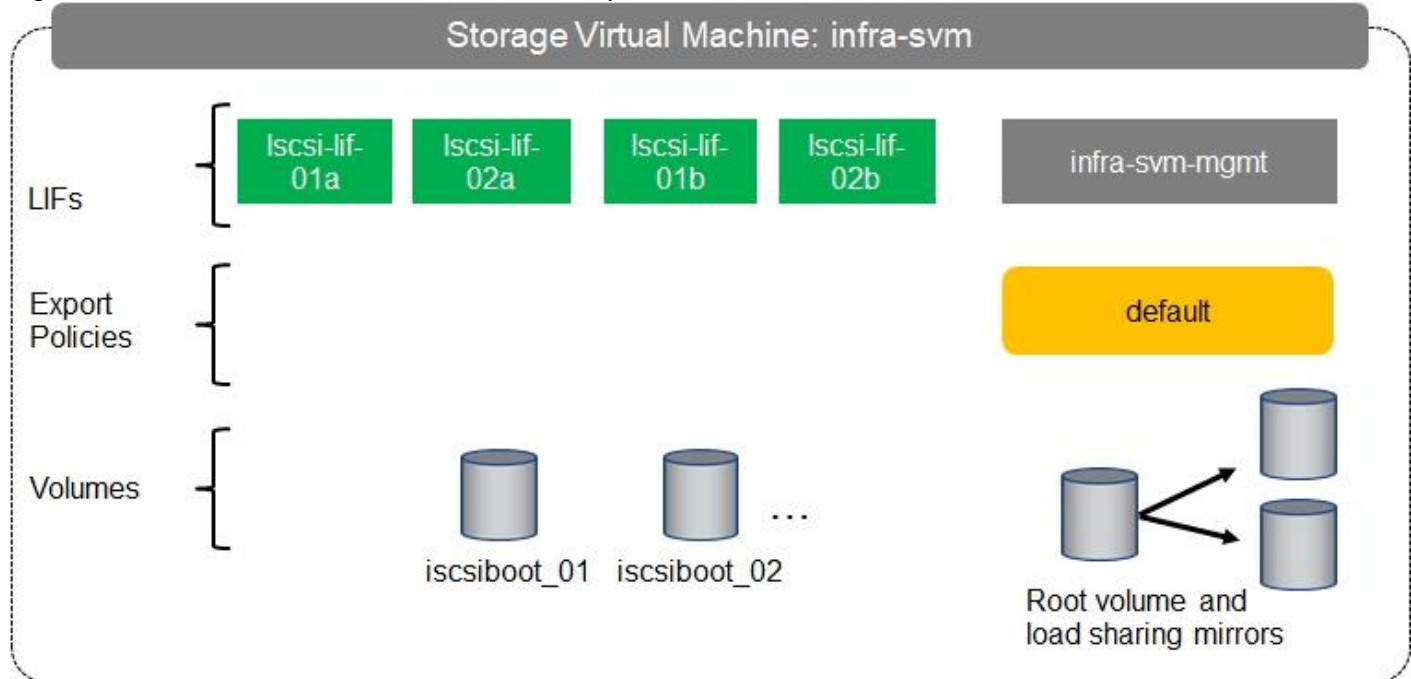


Table 10 ONTAP Software Parameters for Infrastructure SVMs

Cluster Detail	Cluster Detail Value	Value used in CVD setup
Infrastructure SVM management IP	<infra-svm-ip>	192.168.76.36
Infrastructure SVM management IP netmask	<infra-svm-netmask>	255.255.255.0
Infrastructure SVM default gateway	<infra-svm-gateway>	192.168.76.1
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
IQN of Node 1	<server-host-infra-01-iqn>	iqn.2019-09.com.flexpod:hana-node:1
IQN of Node 2	<server-host-infra-02-iqn>	iqn.2019-09.com.flexpod:hana-node:2

Create SVM for the Infrastructure

To create an infrastructure SVM, follow these steps:

1. Run the vserver create command.

```
vserver create -vserver infra-svm -rootvolume infra_rootvol -aggregate aggr2_1 -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 data aggregates to the Infra-svm aggregate list .

```
vserver modify -vserver infra-svm -aggr-list aggr1_1, aggr2_1, aggr1_2, aggr2_2
```

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, follow these 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 aggr2_1 -size 1GB -type DP
volume create -vserver infra-svm -volume infra_rootvol_m02 -aggregate aggr2_2 -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 5min
snapmirror create -source-path infra-svm:infra_rootvol -destination-path
infra-svm:infra_rootvol_m02 -type LS -schedule 5min
```

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, follow these 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 export policy to the infrastructure SVM root volume.

```
volume modify -vserver infra-svm -volume infra_rootvol -policy default
```

Add Infrastructure SVM Management LIF

To add the infrastructure SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver infra-svm -lif infra-svm-mgmt -service-policy default-management -role
data -data-protocol none -home-node <clustername>-02 -home-port e0M -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>
```

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 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. Not required, if SVM creation has been done via ONTAP System Manager

```
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 1000GB -state online -
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 boot LUNs, run the following commands:

```
lun create -vserver infra-svm -volume iscsiboot_01 -lun server-01 -size 80GB -ostype linux -space-reserve disabled
lun create -vserver infra-svm -volume iscsiboot_01 -lun server-02 -size 80GB -ostype linux -space-reserve disabled
```

The example above created boot LUNs for two servers. Repeat the command with a different LUN name to create additional boot LUNs for additional servers

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 IQN information you defined in section Create IQN Pools to create the igroups.

To create igroups, run the following commands:

```
igroup create -vserver Infra-SVM -igroup server-01 -protocol iscsi -ostype linux -initiator <server-host-infra-01-iqn> -portset server_Portset
igroup create -vserver Infra-SVM -igroup server-02 -protocol iscsi -ostype linux -initiator <server-host-infra-02-iqn> -portset server_Portset
```

Repeat this command by using the iqn name of additional servers to create additional igroups for additional servers

Map Boot LUNs to igroups

To map server boot LUNs to igroups, run the following commands:

```
lun map -vserver Infra-SVM -volume iscsiboot_01 -lun server-01 -igroup server-01 -lun-id 0
lun map -vserver Infra-SVM -volume iscsiboot_01 -lun server-02 -igroup server-02 -lun-id 0
```

Repeat this command to map additional boot LUNs to additional servers

Configure SVM for HANA

Table 11 and Figure 6 describe the HANA SVM together with all the required storage objects (volumes, export-policies, and LIFs).

Figure 6 Overview of SAP HANA SVM Components

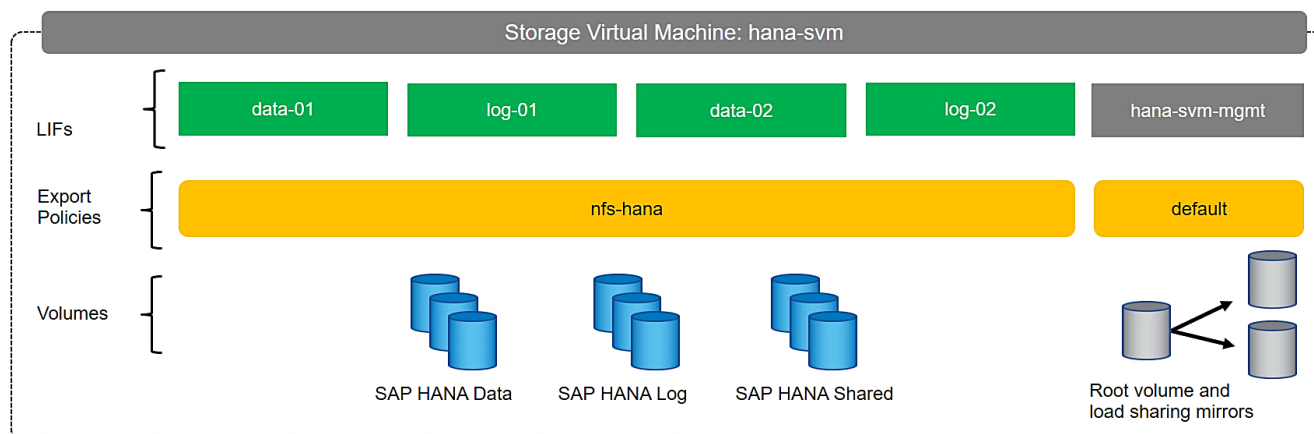


Table 11 ONTAP Software Parameter for HANA SVM

Cluster Detail	Cluster Detail Value	Value used in CVD setup
HANA SVM management IP	<hana-svm-ip>	192.168.76.37
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>	192.168.201.0
NFS Data netmask	<data-netmask>	255.255.255.0
NFS Data LIF node 1 IP	<node01-data_lif01-ip>	192.168.201.11
NFS Data LIF node 2 IP	<node02-data_lif02-ip>	192.168.201.12
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
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, follow these steps:

1. Run the vserver create command.

```
vserver create -vserver hana-svm -rootvolume hana_rootvol -aggregate aggr1_2-rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping NFS.

```
vserver remove-protocols -vserver hana-svm -protocols fcp,cifs,iscsi,nvme
```

3. Add the two data aggregates to the hana-svm aggregate list.

```
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 -v4.1 enabled -v4.1-pnfs enabled -v4.1-read-delegation disabled -v4.1-write-delegation disabled -v4.1-acl disabled -v4-numeric-ids disabled -udp disabled
```

5. Enable a large NFS transfer size.

```
set advanced
vserver nfs modify -vserver hana-svm -tcp-max-transfersize 1048576
set admin
```

6. Set the NFSv4 ID domain.

```
nfs modify -vserver hana-svm -v4-id-domain nfsv4domain.flexpod.com
```

7. Set the NFSv4 lease time.

```
set advanced
nfs modify -vserver hana_svm -v4-lease-seconds 10
set admin
```

8. Set the group ID of the user root to 0.

```
vserver services unix-user modify -vserver hana-svm -user root -primary-gid 0
```

Create Load-Sharing Mirrors

To create a load-sharing mirror of an SVM root volume, follow these 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 aggr2_1 -size 1GB -type DP
volume create -vserver hana-svm -volume hana_rootvol_m02 -aggregate aggr2_2 -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 5min
snapmirror create -source-path hana-svm:hana_rootvol -destination-path hana-svm:hana_rootvol_m02 -type LS -schedule 5min
```

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, follow these 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
```

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 Management Interface and Administrator

To add the HANA SVM administrator and SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver hana-svm -lif hana-svm-mgmt -service-policy default management -role
data -data-protocol none -home-node <clustername>-02 -home-port e0M -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

To create an export policy for the HANA SVM, follow these steps:

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 -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

This section details the preparation of HANA nodes based on SLES 15 and RHEL 7.6.

SAP HANA Node OS Preparation – RHEL for SAP HANA 7.6

This section details the RHEL 7.6 installation and configuration.

OS Installation

To install the OS, follow these steps:



The following procedure shows the RHEL 7.6 installation procedure. Keep the RHEL DVD handy.

This section provides the procedure for Red Hat Enterprise Linux 7.6 Operating System and customizing for SAP HANA requirement.



RHEL 7 must be installed and configured according to the appropriate configuration guide attached to SAP note [2009879](#).

To install the RHEL 7.6 system, follow these steps:

1. Prepare the iSCSI LUN like described in the Storage part of this CVD for the OS.
2. In Cisco UCS Manager, click the Servers tab in the navigation pane.
3. Select Service Profiles > root > HANA-Server01.
4. Click KVM Console.
5. When the KVM Console is launched, click Boot Server.
6. If you using CD click Virtual Media > Activate Virtual Devices.
7. Select Accept this Session for Unencrypted Virtual Media Session then click Apply.
8. Click Virtual Media and Choose Map CD/DVD.
9. Click Browse to navigate ISO media location.
10. Click Map Device.
11. Normally a reboot is necessary to activate this virtual drive.
12. During the reboot the iSCSI targets must be shown. If not check the iSCSI configuration.

```

9  0  SEAGATE   ST300MM0048          NOB1          286102MB
   0  AVAGO     Virtual Drive        RAID1         285148MB

0 JBOD(s) found on the host adapter
1 Virtual Drive(s) found on the host adapter.
Adapter BIOS Disabled. No Logical Drive Handled by BIOS on HA - 0

0 JBOD(s) handled by BIOS
0 Virtual Drive(s) handled by BIOS

Press <Ctrl><R> to Enable BIOS

Cisco VIC iSCSI, Boot Driver Version 4.3(3b)
(C) 2016 Cisco Systems, Inc.
 00:25:b5:0a:00:04 iSCSI NETAPP
 00:25:b5:0a:00:04 iSCSI NETAPP
Option ROM installed successfully

Cisco VIC iSCSI, Boot Driver Version 4.3(3b)
(C) 2016 Cisco Systems, Inc.
 00:25:b5:0b:00:03 iSCSI NETAPP
 00:25:b5:0b:00:03 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 7.6

Install Red Hat Enterprise Linux 7.6
Test this media & install Red Hat Enterprise Linux 7.6

Troubleshooting >

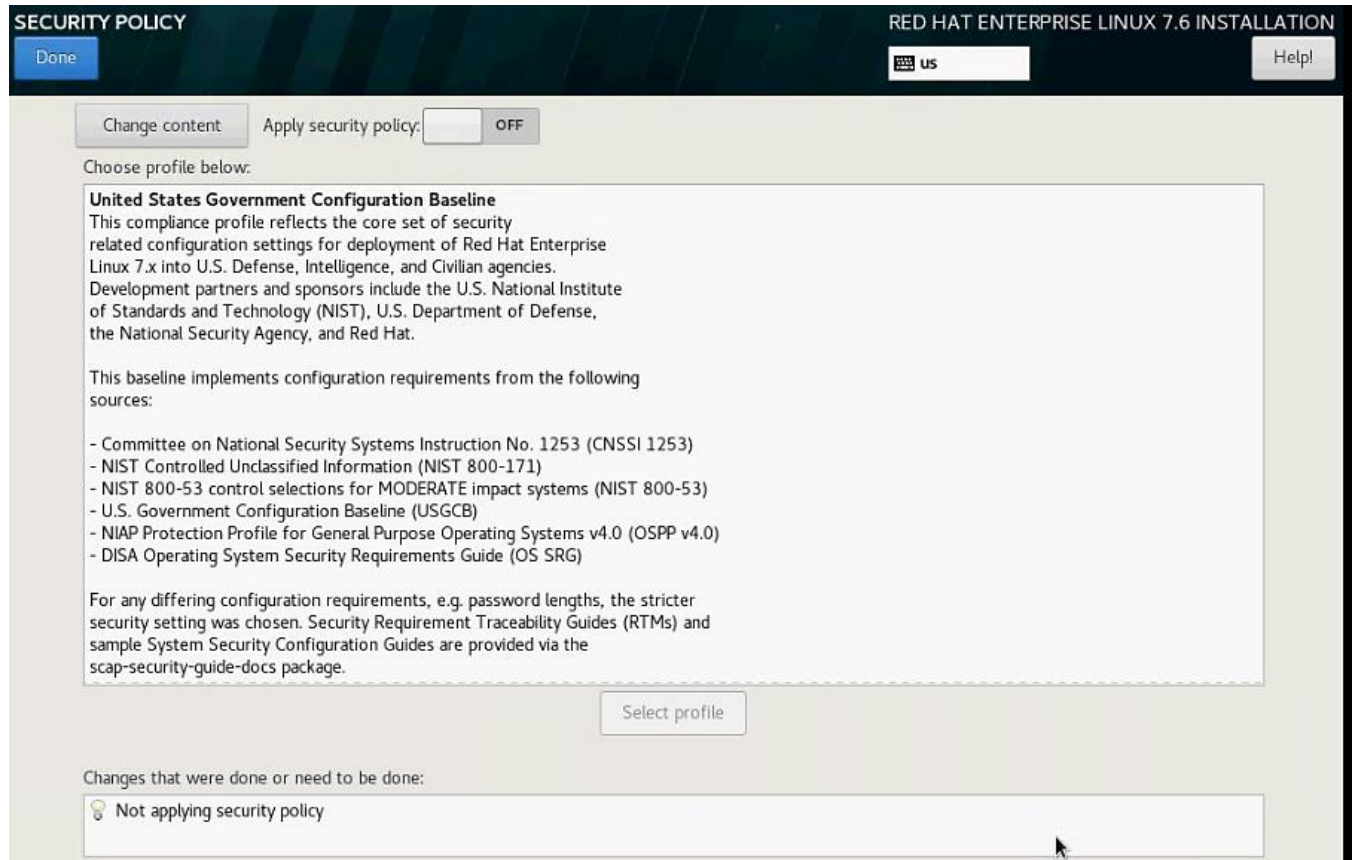
> vmlinuz initrd=initrd.img inst.stage2=hd:LABEL=RHEL-7.6\x20Server.x86_64 quiet rd.iscsi.ibft=1_
```

This is added to make sure the iSCSI targets are discovered appropriately.



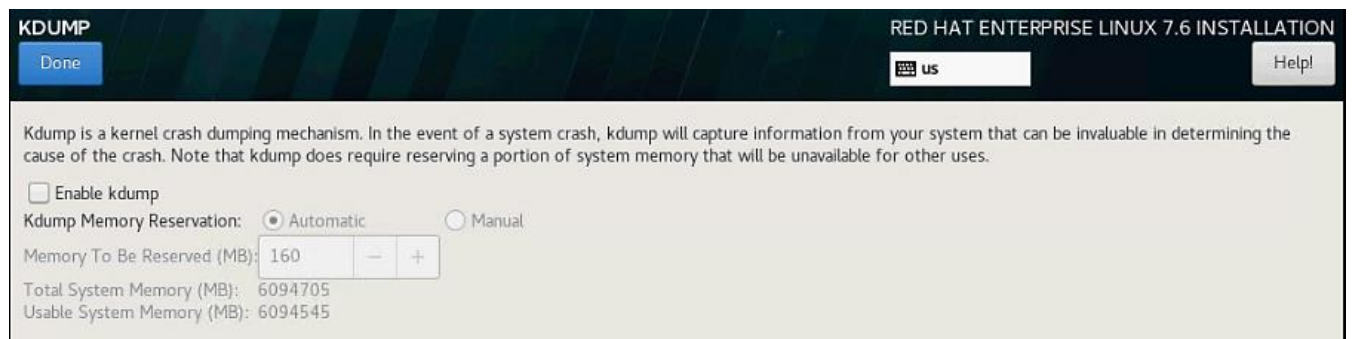
Do not change the system Language (must be English/English).

16. Choose Keyboard and configure your layout.
17. Configure the appropriate Timezone and Time.
18. Click the 'Security Policy' to set the security policy to OFF.

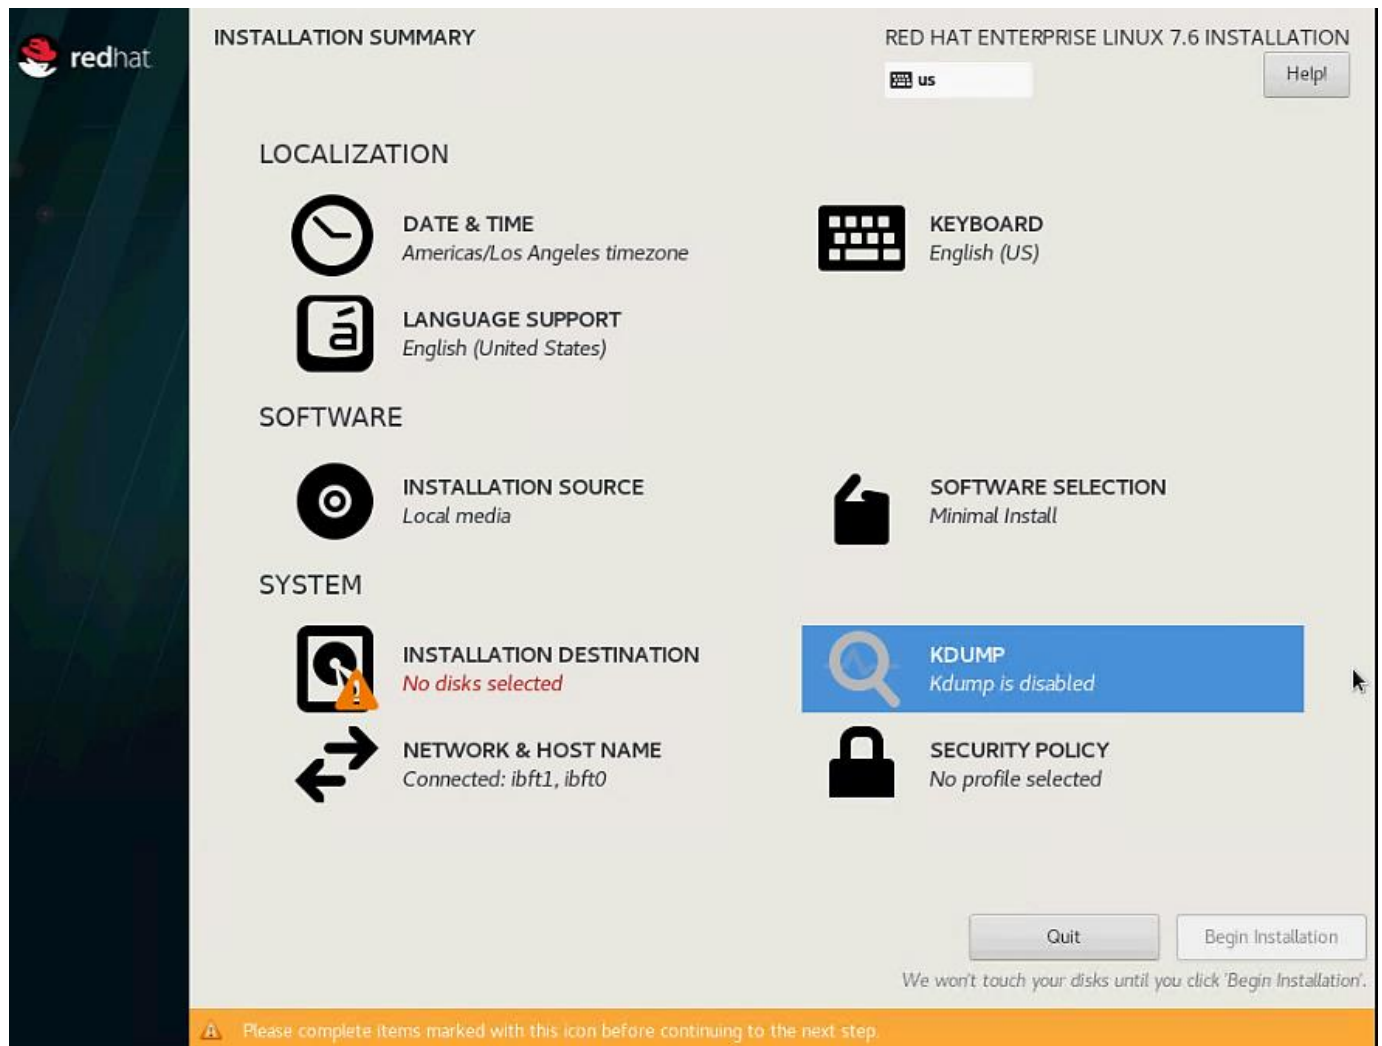


19. Leave the Software section selections as default (Minimal Installation).

20. Disable KDUMP.



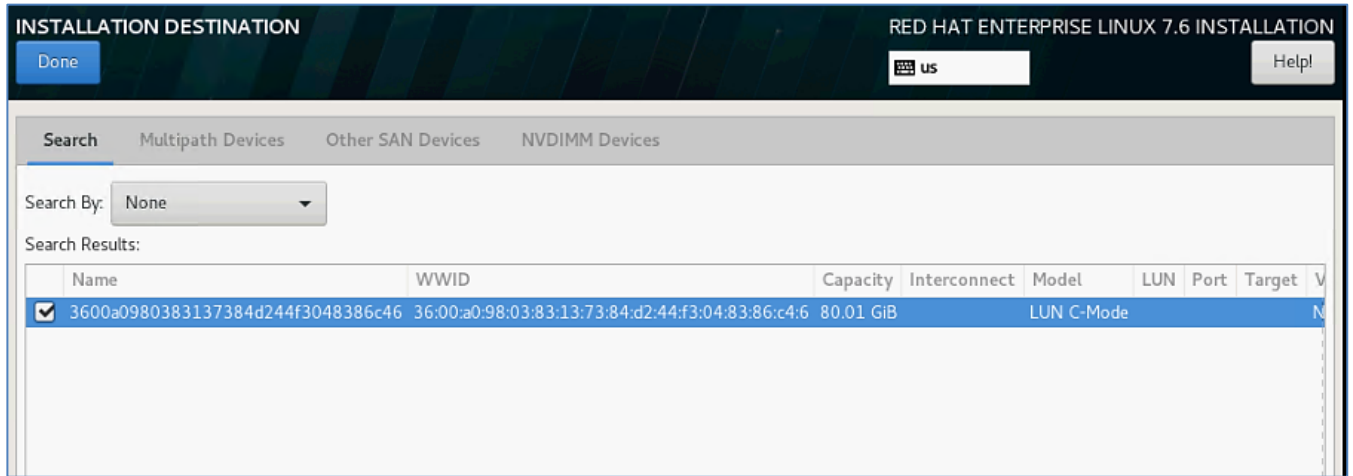
21. Click Installation destination.



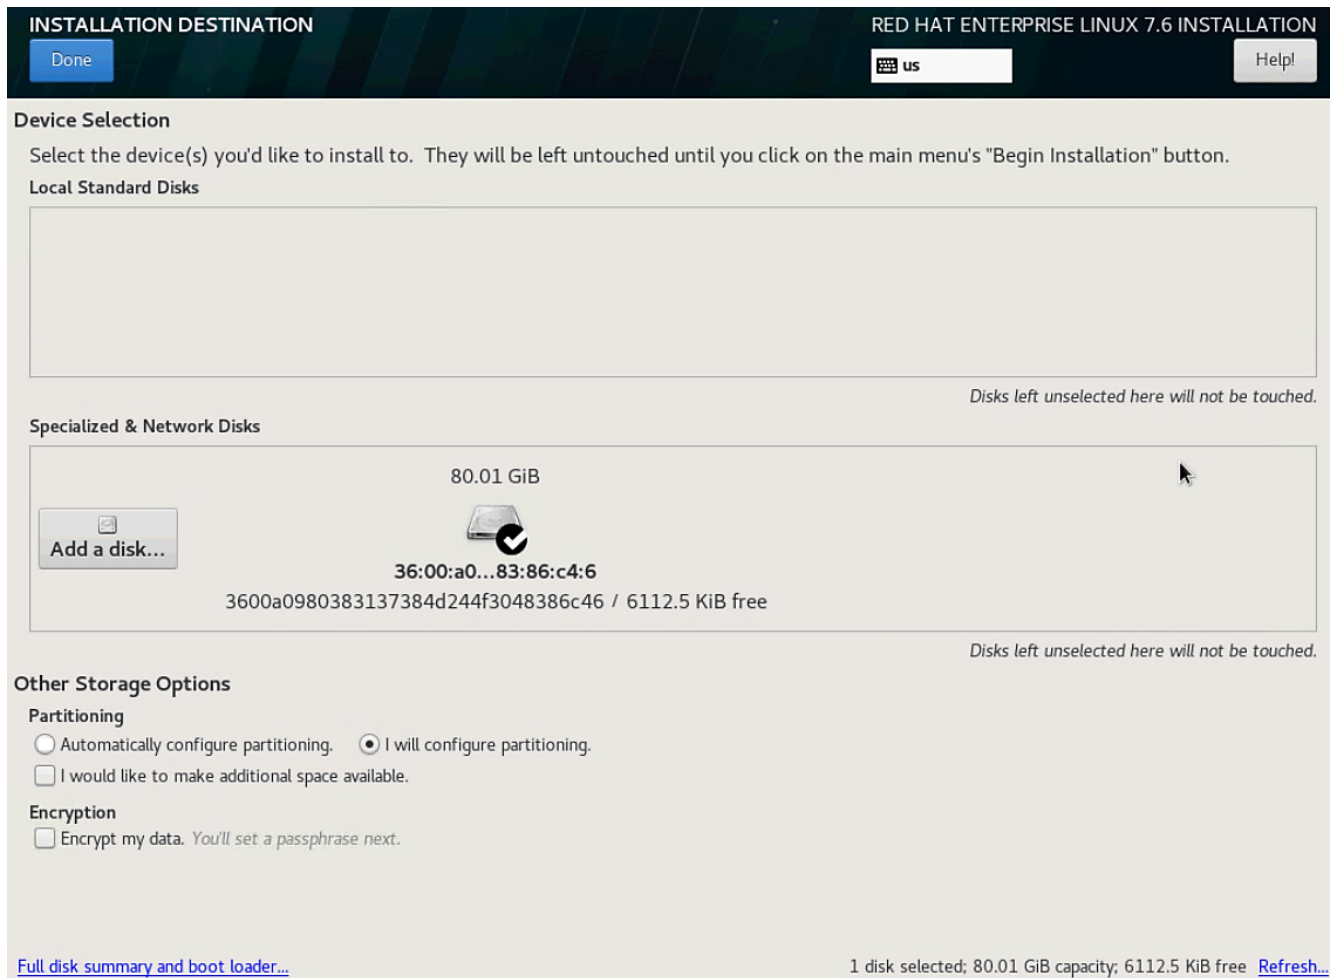
The next screen lists all the local standard disks, if any.

22. Click Specialized and Network Disks section – “Add a Disk”

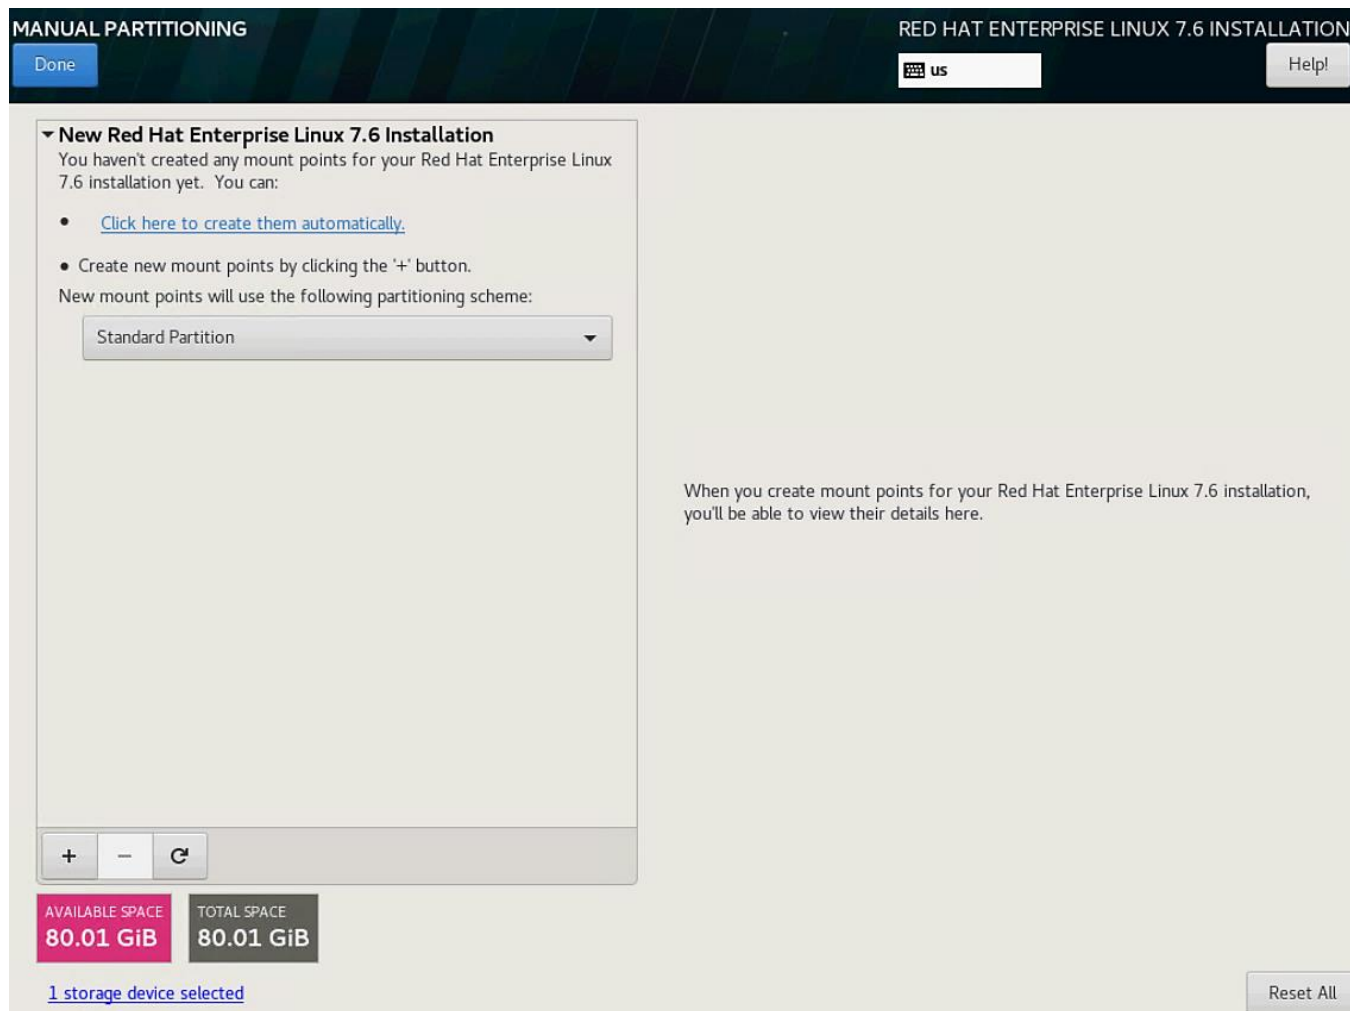
23. Select the discovered iSCSI boot LUN disk and click “Done.”



24. Click Done. Select "I will configure partitioning". Click Done.



25. Select Standard Partition and then select "Click here to create them automatically."



26. Confirm the default Partition table. Click Done. Accept Changes.

MANUAL PARTITIONING RED HAT ENTERPRISE LINUX 7.6 INSTALLATION

Done us Help!

▼ New Red Hat Enterprise Linux 7.6 Installation

SYSTEM	
/boot 3600a0980383137384d244f3048386c46p1	1024 MiB
/ 3600a0980383137384d244f3048386c46p3	75 GiB >
swap 3600a0980383137384d244f3048386c46p2	4096 MiB

+ - ↻

AVAILABLE SPACE 6112.5 KiB **TOTAL SPACE** 80.01 GiB

[1 storage device selected](#)

3600a0980383137384d244f3048386c46p3

Mount Point: /
Desired Capacity: 75 GiB

Device(s):
WWID
36:00:a0:98:03:83:13:73:84:d2:44:f3:04:83:86:c4:6
(3600a0980383137384d244f3048386c46)

Modify...

Device Type: Standard Partition Encrypt

File System: xfs Reformat

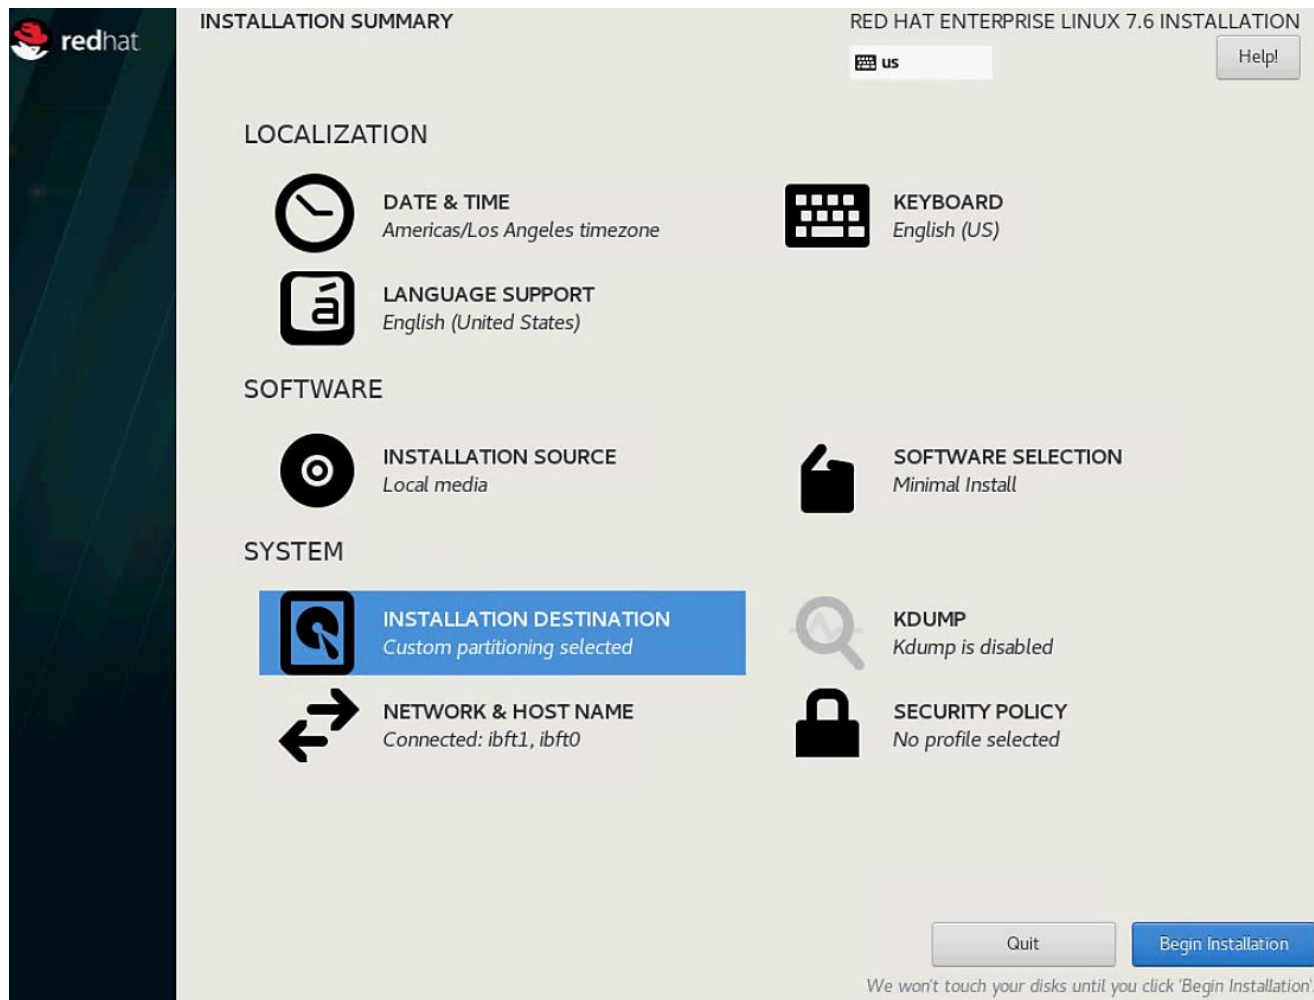
Label: Name: 3600a0980383137384d244f3048386c46

Update Settings

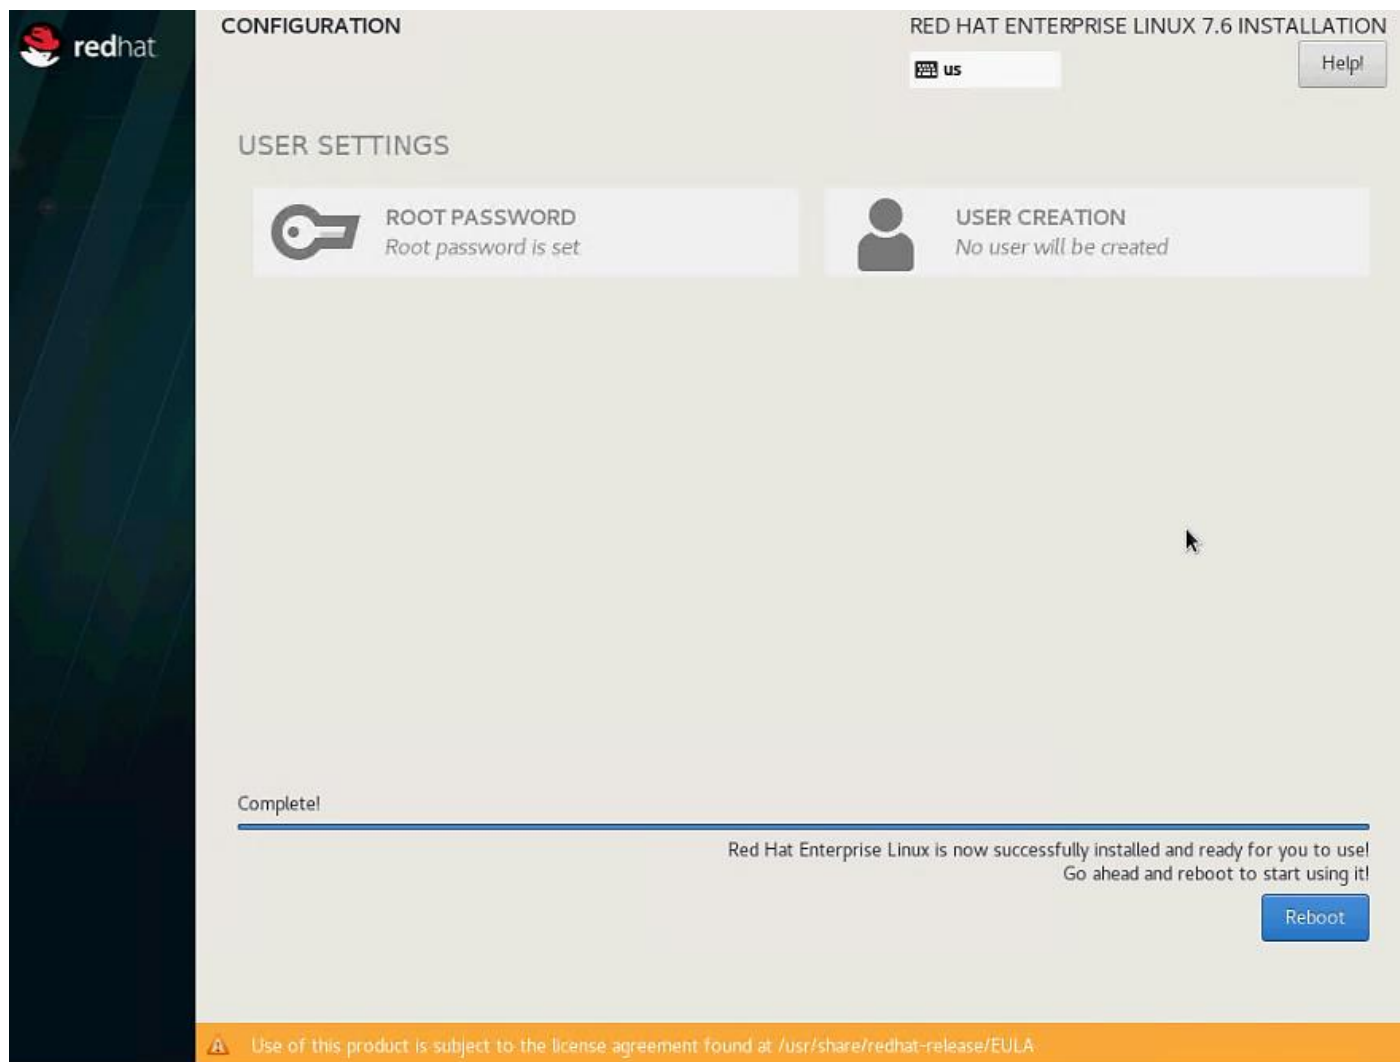
Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.

Reset All

27. Click Begin Installation and then setup the root password.



28. Installation completes. Click Reboot.



Post Installation Tasks

Configure 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.

You could change to traditional device names by setting these Kernel command line parameters `net.ifnames=0 biosdevname=0`. At this time, you can disable IPv6 support with `ipv6.disable=1`.

1. Log into the newly installed system as root.
2. Add the following to the line starting with `GRUB_CMDLINE_LINUX` in `/etc/default/grub`:

```
net.ifnames=0 biosdevname=0 ipv6.disable=1
```

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="iscsi_firmware ip=ibft rhgb net.ifnames=0 biosdevname=0 ipv6.disable=1"
GRUB_DISABLE_RECOVERY="true"
~
.
```

3. Rebuild the /boot/grub2/grub.cfg file.

```
#grub2-mkconfig -o /boot/grub2/grub.cfg
```

```
[root@localhost ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-3.10.0-957.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-957.el7.x86_64.img
Found linux image: /boot/vmlinuz-0-rescue-b1b1094271dd4dac833c74a722530848
Found initrd image: /boot/initramfs-0-rescue-b1b1094271dd4dac833c74a722530848.img
Found Red Hat Enterprise Linux Server release 7.6 (Maipo) on /dev/sde2
done
[root@localhost ~]#
```

4. Reboot for effect the disabling of consistent network device naming. Log back in and check the interfaces.

```
Red Hat Enterprise Linux Server 7.6 (Maipo)
Kernel 3.10.0-957.el7.x86_64 on an x86_64

localhost login: root
Password:
Last login: Tue Oct 15 07:54:07 from 192.168.76.20
[root@localhost ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    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: ibft0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0a:00:04 brd ff:ff:ff:ff:ff:ff
    inet 192.168.128.51/24 brd 192.168.128.255 scope global ibft0
        valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0a:00:05 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0b:00:04 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0a:00:07 brd ff:ff:ff:ff:ff:ff
6: ibft1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0b:00:03 brd ff:ff:ff:ff:ff:ff
    inet 192.168.129.51/24 brd 192.168.129.255 scope global ibft1
        valid_lft forever preferred_lft forever
7: eth5: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0a:00:06 brd ff:ff:ff:ff:ff:ff
8: eth6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0b:00:05 brd ff:ff:ff:ff:ff:ff
[root@localhost ~]#
```



The iSCSI vNICs appear as ibft0 with ifcfg-eth0 interface config file and other iSCSI vNIC as ibft1 and correspondingly ifcfg-eth4 config file. We will have to create the ifcfg-eth1, ifcfg-eth2, ifcfg-eth3, ifcfg-eth5 and ifcfg-eth6 files.

- Configure the admin/management network, hostname and default GW. Identify the Admin/Mgmt vNIC using the MAC address to eth interface mapping at the OS level with “ip address” command [as above] and on the Network tab of service profile under the vNICs.

vNICs	
Name	MAC Address
vNIC HANA-AppServer	00:25:B5:0A:00:07
vNIC HANA-data	00:25:B5:0A:00:06
vNIC HANA-log	00:25:B5:0B:00:05
vNIC HANA-node-admin	00:25:B5:0A:00:05
vNIC HANA-node-backup	00:25:B5:0B:00:04

- Create the configuration file for the interface as below. Default Route and Gateway setting are done at the interface level.

```
#cd /etc/sysconfig/network-scripts

#vi ifcfg-eth1
TYPE=Ethernet
BOOTPROTO=static
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
NAME=admin
DEVICE=eth1
ONBOOT=yes
IPADDR=192.168.76.202
PREFIX=24
GATEWAY=192.168.76.1

#vi ifcfg-eth2
TYPE=Ethernet
BOOTPROTO=static
DEFROUTE=no
IPV4_FAILURE_FATAL=no
NAME=backup
DEVICE=eth2
ONBOOT=yes
IPADDR=192.168.224.202
PREFIX=24

#vi ifcfg-eth3
TYPE=Ethernet
```

```

BOOTPROTO=static
DEFROUTE=no
IPV4_FAILURE_FATAL=no
NAME=appserver
DEVICE=eth3
ONBOOT=yes
IPADDR=192.168.223.202
PREFIX=24

#vi ifcfg-eth5
TYPE=Ethernet
BOOTPROTO=static
DEFROUTE=no
IPV4_FAILURE_FATAL=no
NAME=hana-data
DEVICE=eth5
ONBOOT=yes
IPADDR=192.168.201.202
PREFIX=24

#vi ifcfg-eth6
TYPE=Ethernet
BOOTPROTO=static
DEFROUTE=no
IPV4_FAILURE_FATAL=no
NAME=hana-log
DEVICE=eth6
ONBOOT=yes
IPADDR=192.168.228.202
PREFIX=24

```

7. Update the hostname with hostnamectl command

```
#hostnamectl set-hostname <hana-node-hostname>
```

8. Restart the Network.

```
#systemctl restart network
```

Update the Red Hat System

Refer to the Red Hat KB article <https://access.redhat.com/solutions/65300> to prepare the system to access the Red Hat Subscription Manager through proxy.

To register the system to Red Hat portal and attach SAP HANA subscription, follow the instructions in the Red Hat KB articles:

<https://access.redhat.com/solutions/2318061> and <https://access.redhat.com/solutions/2334521>

To update the Red Hat System, follow these steps:

1. Set the release to the right minor release and clear the yum cache and subscribe to the right channels.

```
subscription-manager register
subscription-manager list --available --all
```

```

subscription-manager attach --pool=<<Pool-ID>>

subscription-manager release --set=7.6
yum clean all

subscription-manager repos --enable="rhel-sap-hana-for-rhel-7-server-rpms"
subscription-manager repos --enable="rhel-ha-for-rhel-7-server-rpms"
subscription-manager repos --enable="rhel-sap-hana-for-rhel-7-server-eus-rpms" --
enable="rhel-7-server-eus-rpms"
subscription-manager repos --enable="rhel-ha-for-rhel-7-server-eus-rpms"

```

2. Check for the available repositories.

```

[root@fprhel01 ~]# yum repolist
rhel-7-server-eus-rpms
| 3.5 kB 00:00:00
rhel-7-server-rpms
| 3.5 kB 00:00:00
rhel-ha-for-rhel-7-server-eus-rpms
| 3.4 kB 00:00:00
rhel-ha-for-rhel-7-server-rpms
| 3.4 kB 00:00:00
rhel-sap-hana-for-rhel-7-server-eus-rpms
| 4.0 kB 00:00:00
rhel-sap-hana-for-rhel-7-server-rpms
| 4.0 kB 00:00:00
(1/18): rhel-7-server-eus-rpms/x86_64/group
| 772 kB 00:00:00
(2/18): rhel-7-server-eus-rpms/x86_64/updateinfo
| 3.0 MB 00:00:00
(3/18): rhel-7-server-rpms/x86_64/group
| 773 kB 00:00:00
(4/18): rhel-7-server-rpms/x86_64/updateinfo
| 3.2 MB 00:00:00
(5/18): rhel-ha-for-rhel-7-server-eus-rpms/x86_64/group
| 14 kB 00:00:00
(6/18): rhel-ha-for-rhel-7-server-eus-rpms/x86_64/updateinfo
| 106 kB 00:00:00
(7/18): rhel-ha-for-rhel-7-server-rpms/x86_64/group
| 16 kB 00:00:00
(8/18): rhel-ha-for-rhel-7-server-eus-rpms/x86_64/primary_db
| 377 kB 00:00:00
(9/18): rhel-ha-for-rhel-7-server-rpms/x86_64/updateinfo
| 113 kB 00:00:00
(10/18): rhel-7-server-eus-rpms/x86_64/primary_db
| 62 MB 00:00:01
(11/18): rhel-sap-hana-for-rhel-7-server-eus-rpms/x86_64/group
| 1.6 kB 00:00:00
(12/18): rhel-ha-for-rhel-7-server-rpms/x86_64/primary_db
| 376 kB 00:00:00
(13/18): rhel-sap-hana-for-rhel-7-server-eus-rpms/x86_64/updateinfo
| 25 kB 00:00:00
(14/18): rhel-sap-hana-for-rhel-7-server-rpms/x86_64/group
| 3.7 kB 00:00:00
(15/18): rhel-sap-hana-for-rhel-7-server-eus-rpms/x86_64/primary_db
| 17 kB 00:00:00

```

```
(16/18): rhel-sap-hana-for-rhel-7-server-rpms/x86_64/updateinfo
| 26 kB 00:00:00
(17/18): rhel-7-server-rpms/x86_64/primary_db
| 58 MB 00:00:02
(18/18): rhel-sap-hana-for-rhel-7-server-rpms/x86_64/primary_db
| 16 kB 00:00:00
repo id
repo name
status
rhel-7-server-eus-rpms/x86_64
Red Hat Enterprise Linux 7 Server - Extended Update Support (RPMs)
24,740
rhel-7-server-rpms/x86_64
Red Hat Enterprise Linux 7 Server (RPMs)
24,539
rhel-ha-for-rhel-7-server-eus-rpms/x86_64
Red Hat Enterprise Linux High Availability (for RHEL 7 Server) - Extended Update
Support (RPMs) 635
rhel-ha-for-rhel-7-server-rpms/x86_64
Red Hat Enterprise Linux High Availability (for RHEL 7 Server) (RPMs)
632
rhel-sap-hana-for-rhel-7-server-eus-rpms/x86_64
RHEL for SAP HANA (for RHEL 7 Server) Extended Update Support (RPMs)
55
rhel-sap-hana-for-rhel-7-server-rpms/x86_64
Red Hat Enterprise Linux for SAP HANA (RHEL 7 Server) (RPMs)
53
repolist: 50,654
[root@fprhel01 ~]#
```

3. Install the base package group:

```
yum -y groupinstall base
```

4. Install other additional required packages required for running SAP HANA on RHEL 7:

```
yum install gtk2 libicu xulrunner sudo tcsh libssh2 expect cairo graphviz iptraf-ng
krb5-workstation krb5-libs libpng12 nfs-utils lm_sensors rsyslog openssl
PackageKit-gtk3-module libcanberra-gtk2 libtool-ltdl xorg-x11-xauth numactl
xfsprogs net-tools bind-utils3.2.1 compat-sap-c++-7 libatomic compat-sap-c++-6
compat-sap-c++-6 ntp ntpdate
```

5. Disable SELinux:

```
sed -i 's/\(SELINUX=enforcing\|SELINUX=permissive\)/SELINUX=disabled/g' \
/etc/selinux/config /etc/sysconfig/selinux
```

6. Disabling the firewall:

```
systemctl stop firewalld
systemctl disable firewalld
```

7. Edit the file /etc/ntp.conf to reflect the appropriate ntp servers for the region and start the ntp service:

```
systemctl enable ntpd.service
```

```
systemctl start ntpd.service
systemctl restart systemd-timedated.service
```

8. Starting with Red Hat Enterprise Linux 7.2 the kernel parameters must be set in a configuration file in the /etc/sysctl.d directory. For example, a configuration file with the name 91-NetApp-HANA.conf must be created.

```
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.ipv4.tcp_rmem = 65536 16777216 16777216
net.ipv4.tcp_wmem = 65536 16777216 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

9. Disable ABRT and Core dumps:

```
systemctl disable abrt-d
systemctl disable abrt-ccpp
systemctl stop abrt-d
systemctl stop abrt-ccpp
```

10. 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
```

11. Increase the maximum number of processes a user is allowed to run. For the same, create the file /etc/security/limits.d/99-sapsys.conf with the following content

```
@sapsys soft nproc unlimited
@sapsys hard nproc unlimited
```

12. Add symbolic links - Since SAP HANA is built on a different Linux Distribution, some of the library names used during the build process doesn't match with the library names used on RHEL7.

```
ln -s /usr/lib64/libssl.so.10 /usr/lib64/libssl.so.1.0.1
ln -s /usr/lib64/libcrypto.so.10 /usr/lib64/libcrypto.so.1.0.13.10
```

13. Reboot the system to effect kernel switch and SELinux settings.

Implement SAP Notes Recommendations

To configure optimal settings for running HANA or HANA2 on RHEL for SAP HANA 7.6, follow the instructions in the [SAP Note 2292690](#).

To implement the SAP notes recommendations, follow these steps:

1. The tuned profile "sap-hana", which is provided by Red Hat as part of RHEL 7 for SAP HANA, contains many of the settings mentioned in the SAP note and also configures some additional settings. Therefore the "sap-hana" tuned profile must be activated on all systems running SAP HANA.

- a. Install tuned-profiles-sap-hana

```
#yum -y install tuned-profiles-sap-hana
[root@fprhel01 ~]# yum -y install tuned-profiles-sap-hana
Loaded plugins: langpacks, product-id, search-disabled-repos, subscription-manager, versionlock
Package tuned-profiles-sap-hana-2.11.0-5.el7_7.1.noarch already installed and latest version
Nothing to do
[root@fprhel01 ~]# █
```

- b. Start and enable the tuned:

```
#systemctl start tuned
#systemctl enable tuned
```

- c. Apply the profile for sap-hana

```
#tuned-adm profile sap-hana
```

- d. Verify the solution applied

```
#tuned-adm verify
[root@fprhel01 ~]# tuned-adm verify
Verification succeeded, current system settings match the preset profile.
See tuned log file ('/var/log/tuned/tuned.log') for details.
[root@fprhel01 ~]# █
```

Install Cisco VIC Drivers

To download the Cisco UCS Drivers ISO bundle, which contains most Cisco UCS Virtual Interface Card drivers, follow these 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 (Cisco 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 4.0(4d).

8. Click 4.0(4d) Version.

9. Download ISO image of Cisco UCS-related drivers `ucs-bxxx-drivers-linux.4.0.4`.
10. Choose your download method and follow the prompts to complete your driver download.
11. Browse the downloaded drivers iso -> `ucs-bxxx-drivers-linux.4.0.4\Network\Cisco\VIC\RHEL\RHEL7.6` and copy `kmod-enic-3.2.210.18-738.12.rhel7u6.x86_64.rpm` to `/opt` of the HANA node.
12. Update the enic driver:

```
rpm -ivh /opt/kmod-enic-3.2.210.18-738.12.rhel7u6.x86_64.rpm
```

13. Browse the downloaded drivers iso -> `ucs-bxxx-drivers-linux.4.0.4\Storage\Cisco\VIC\RHEL\RHEL7.6` and copy `kmod-fnic-2.0.0.42-77.0.rhel7u6.x86_64.rpm` to also `/opt` of the HANA node.
14. Update the enic driver:

```
rpm -ivh /opt/ kmod-fnic-2.0.0.42-77.0.rhel7u6.x86_64.rpm
```

SAP HANA Node OS Preparation – SLES for SAP 15

In this section, SLES for SAP 15 installation and configuration is detailed.

OS Installation

To install the OS, follow these steps:



The following steps show the SLES 15 installation procedure. Keep the SLES 15 DVD handy.

1. Refer the [2578899 – SUSE Linux Enterprise Server 15: Installation Note](#) for installation instructions.
2. On the UCSM page, Servers -> Service Profiles -> root -> Sub-Organizations -> HANA – Right-click HANA-node01 and select KVM console.
3. After the KVM console is launched, click Boot Server.
4. Choose Virtual Media > Activate Virtual Devices.
 - a. For Unencrypted Virtual Media Session, select Accept this Session and then click Apply.
 - b. Click Virtual Media and choose Map CD/DVD.
 - c. Click Browse to navigate to the ISO media location. Select `SLE-15-Installer-DVD-x86_64-GM-DVD1.iso`. Click Open.
 - d. Click Map Device.
5. At server boot time, during the check the of VIC FC boot driver version, it recognizes the NetApp iSCSI target via four paths; two each per VIC. This verifies the server to storage connectivity.

```
2 JBOD(s) found on the host adapter
0 Virtual Drive(s) found on the host adapter.
Adapter BIOS Disabled. No Logical Drive Handled by BIOS on HA - 0

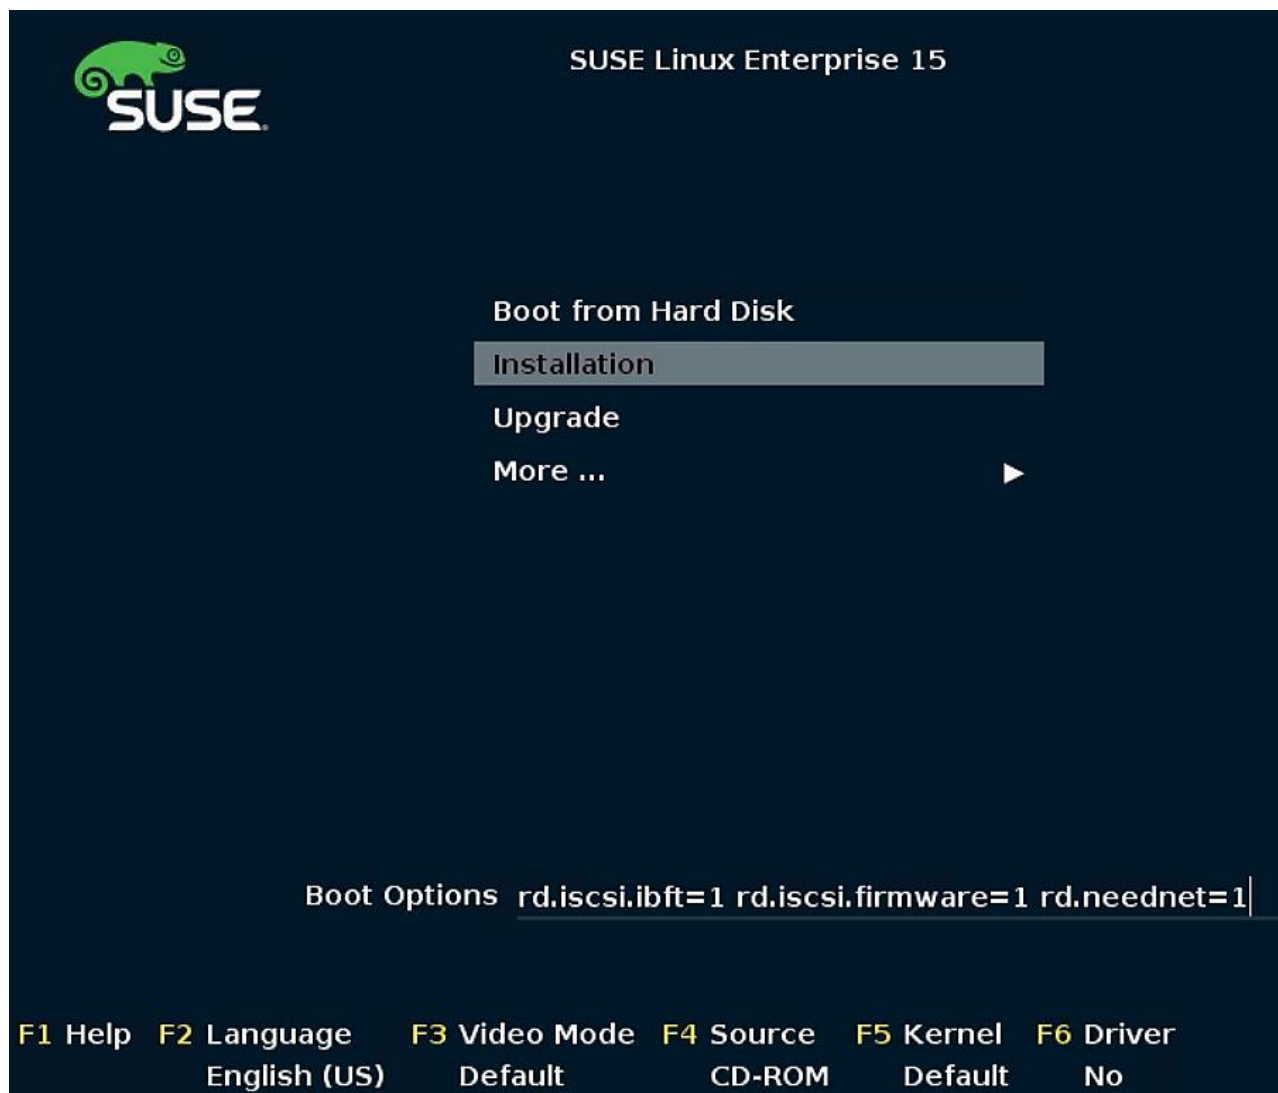
0 JBOD(s) handled by BIOS
0 Virtual Drive(s) handled by BIOS

Press <Ctrl><R> to Enable BIOS
```

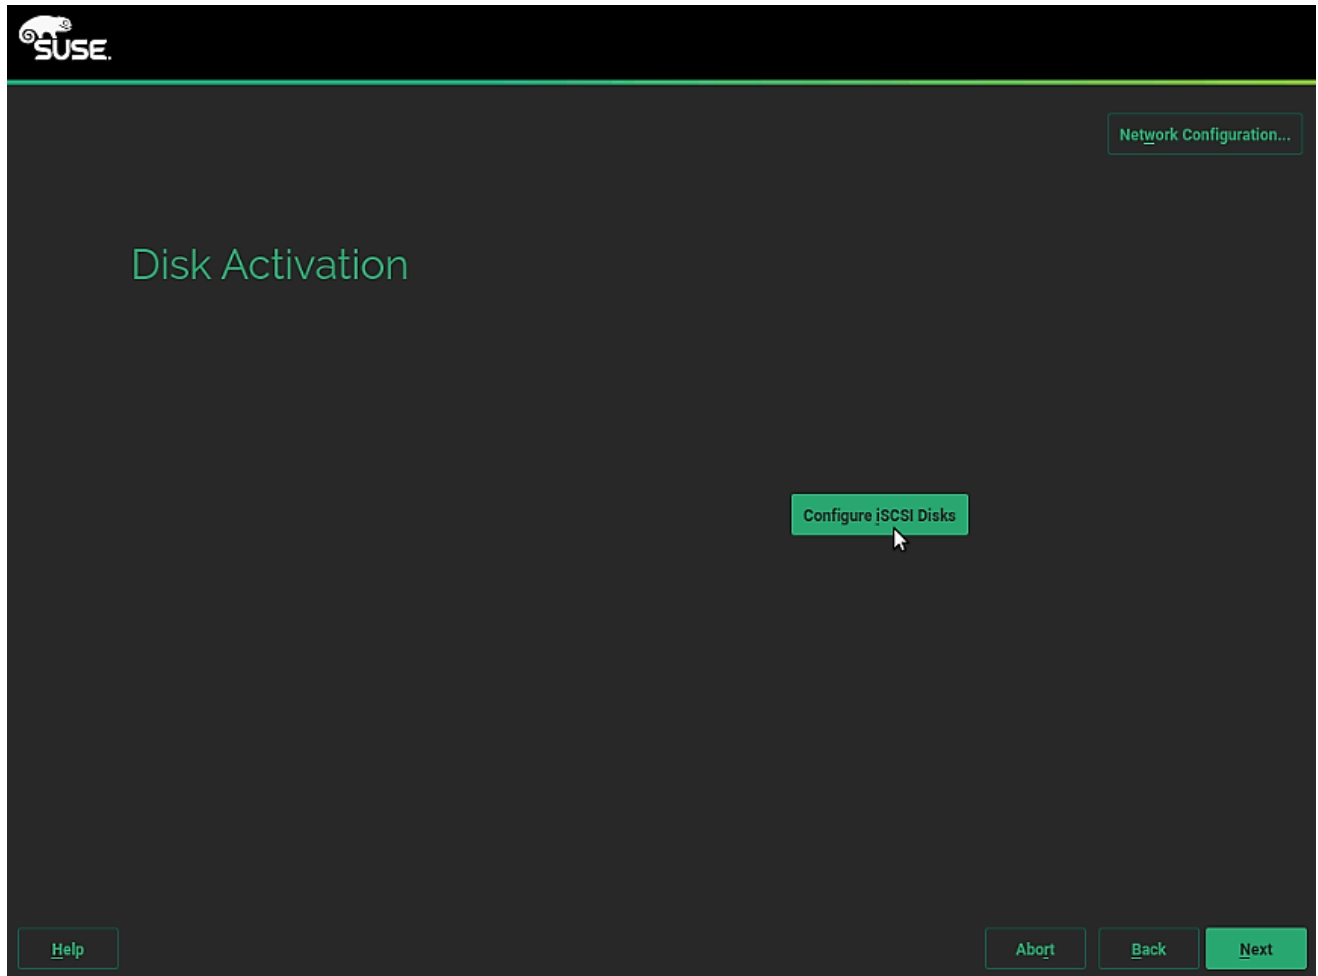
```
Cisco VIC iSCSI, Boot Driver Version 5.0(3c)
(C) 2016 Cisco Systems, Inc.
 00:25:b5:0a:00:0a iSCSI NETAPP
 00:25:b5:0a:00:0a iSCSI NETAPP
Option ROM installed successfully

Cisco VIC iSCSI, Boot Driver Version 5.0(3c)
(C) 2016 Cisco Systems, Inc.
 00:25:b5:0b:00:08 iSCSI NETAPP
 00:25:b5:0b:00:08 iSCSI NETAPP
Option ROM installed successfully
```

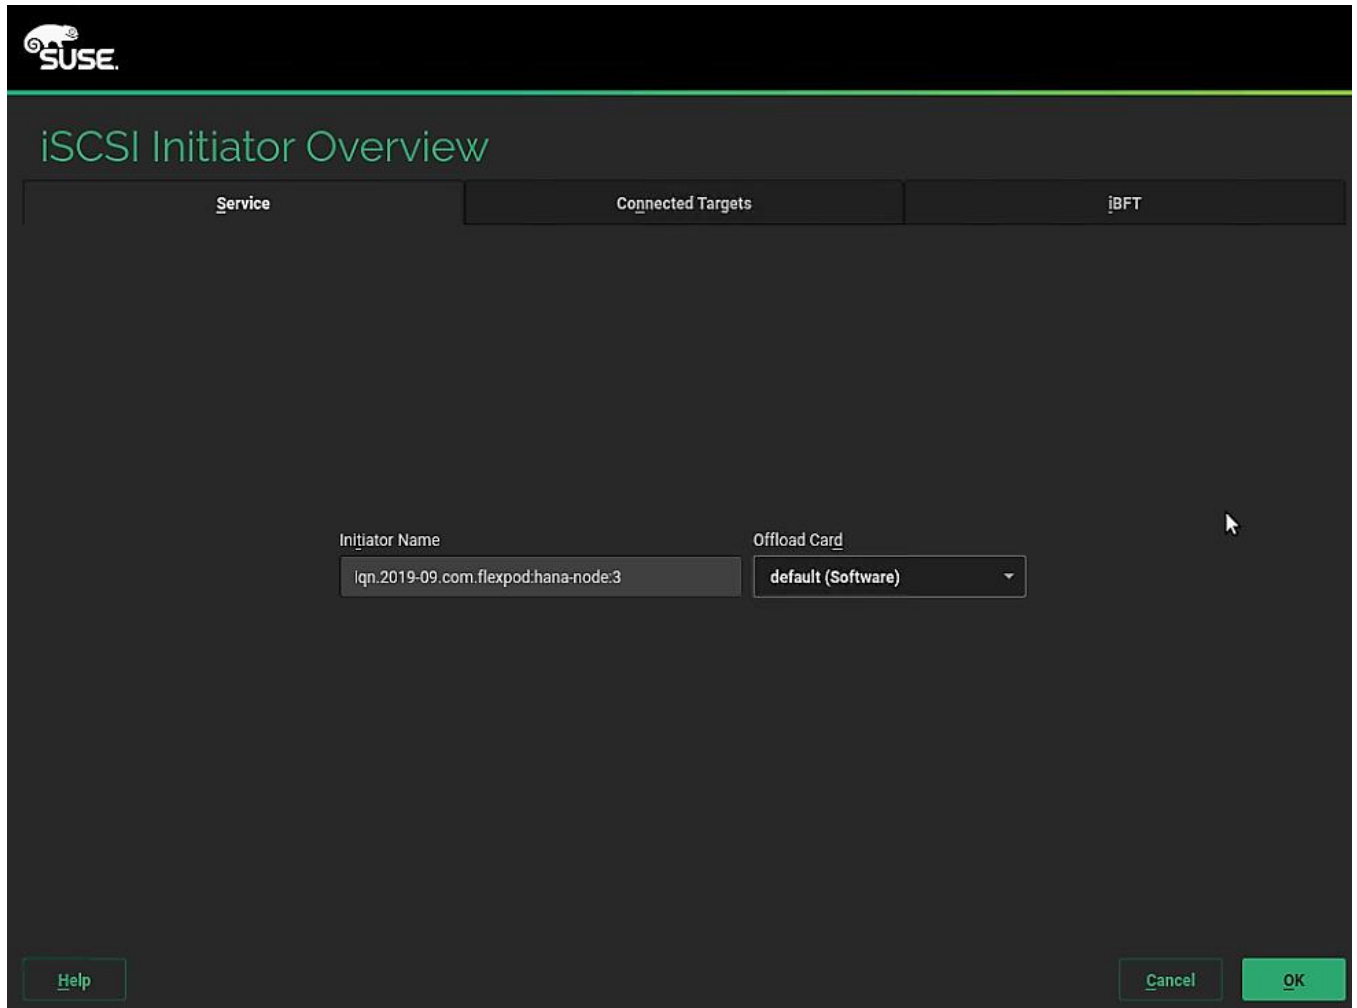
- The System will automatically boot from the ISO image. Scroll down on the selection menu to Installation option and key in the boot options- rd.iscsi.ibft=1 (Turn on iBFT autoconfiguration for the interfaces), rd.iscsi.firmware=1 (to read the iscsi parameter from the BIOS firmware) and rd.neednet=1 (bring up network even without netroot set). Click Return.



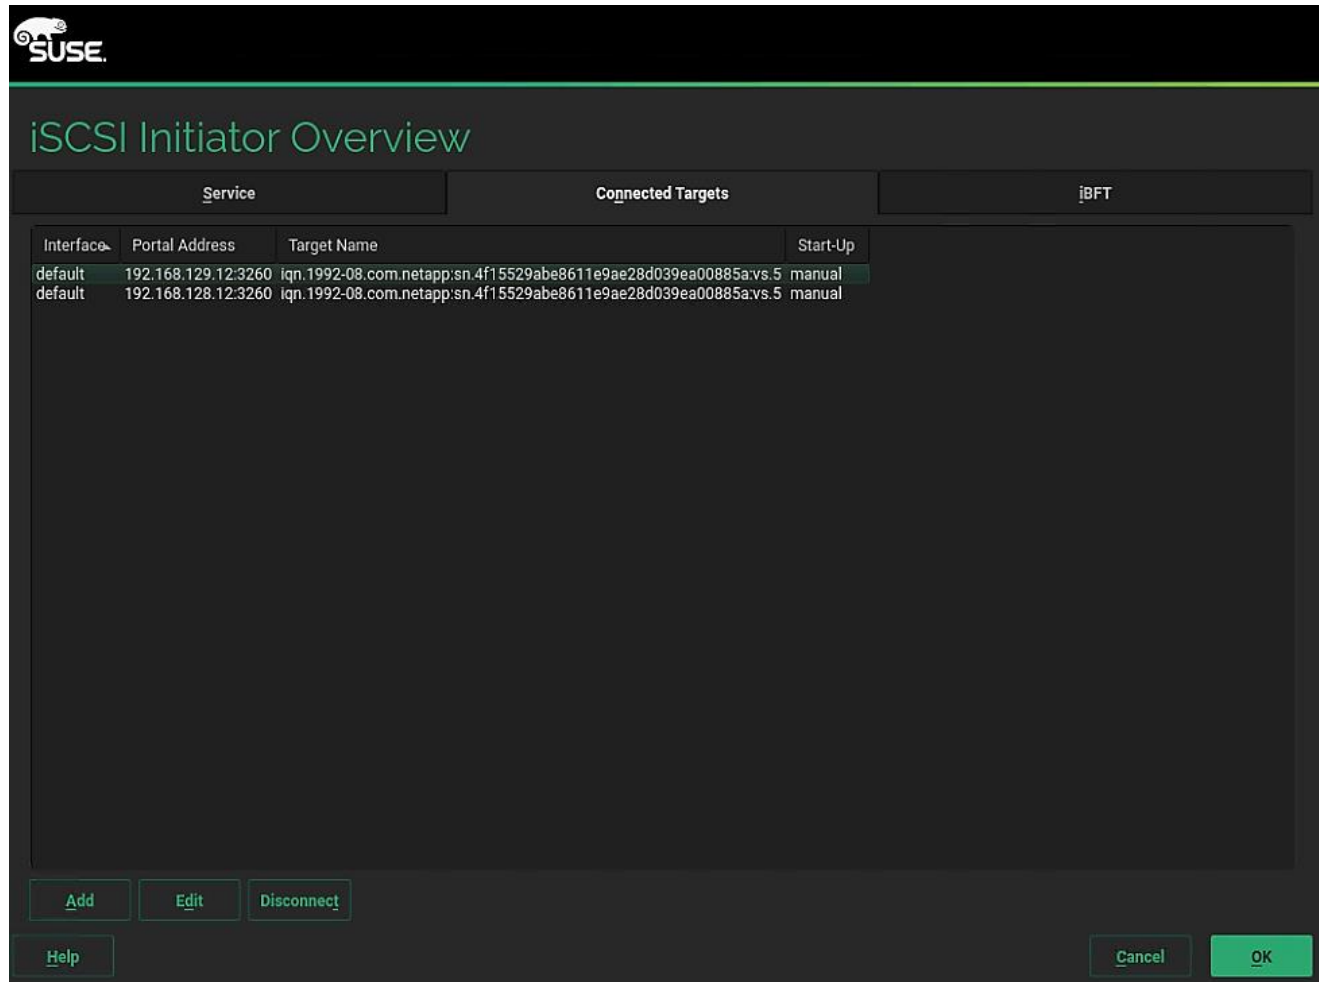
7. On the first “Language, Keyboard and License Agreement” page, select the Language of choice and Keyboard Layout. Select “SUSE Linux Enterprise Server for SAP Applications 15” for Product to Install and click Next.
8. SUSE Linux Enterprise Server for SAP Applications 15 License Agreement – Select “I Agree to the License Terms” . Click Next.
9. *Disk Activation* – Click Configure iSCSI Disks.




10. On the iSCSI Initiator tab, ensure the automatically picked up Initiator Name is correct. [If not update the Initiator Name defined in the service profile.]



11. Click the Connected Targets tab, verify the interface can login to the target via all 4 ports.



12. If not by default, click Add perform iSCSI Initiator Discovery. Key in the IP address of the remaining iSCSI_LIF Ip Address. Click Next.



iSCSI Initiator Discovery

IP Address

Port

No Discovery Authentication

Authentication by Initiator

Username

Password

Authentication by Targets

Username

Password


Help

Cancel

Back

Next

13. The remaining paths are discovered and added. Click Connect selecting newly added paths one at a time [with startup value False]



iSCSI Initiator Discovery

Interface	Portal Address	Target Name	Start-Up
default	192.168.128.11:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	False
default	192.168.129.12:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	True
default	192.168.129.11:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	False
default	192.168.128.12:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	True

Warning

The target with this TargetName is already connected. Make sure that multipathing is enabled to prevent data corruption.

[Continue](#) [Cancel](#)

[Connect](#)

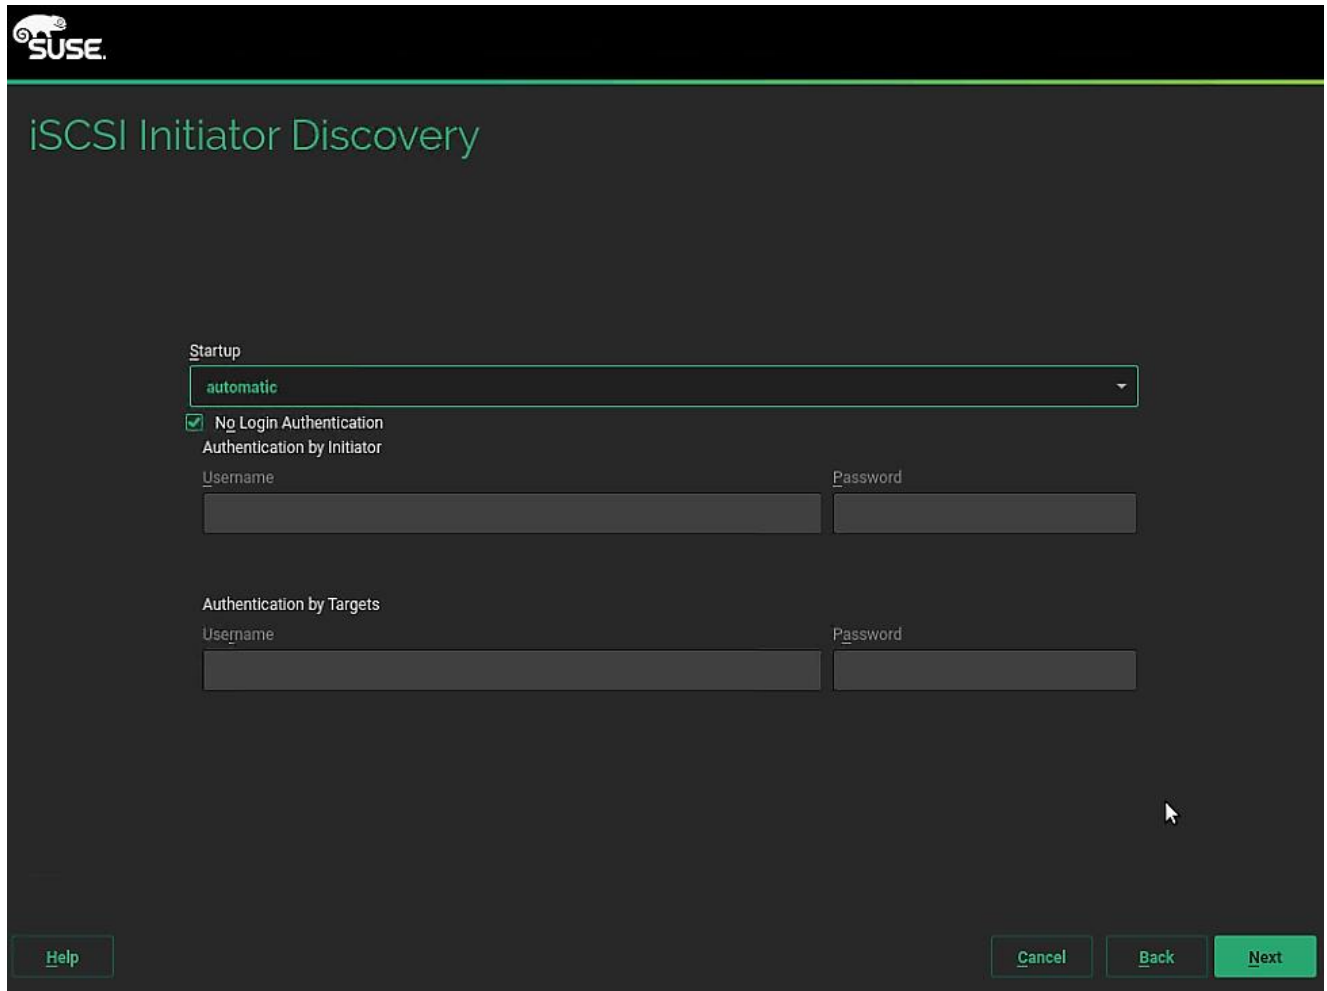
[Help](#)

[Cancel](#)

[Back](#)

[Next](#)

14. Select Continue for the Warning. Set Startup to “automatic.”



15. Do the same for remaining path.

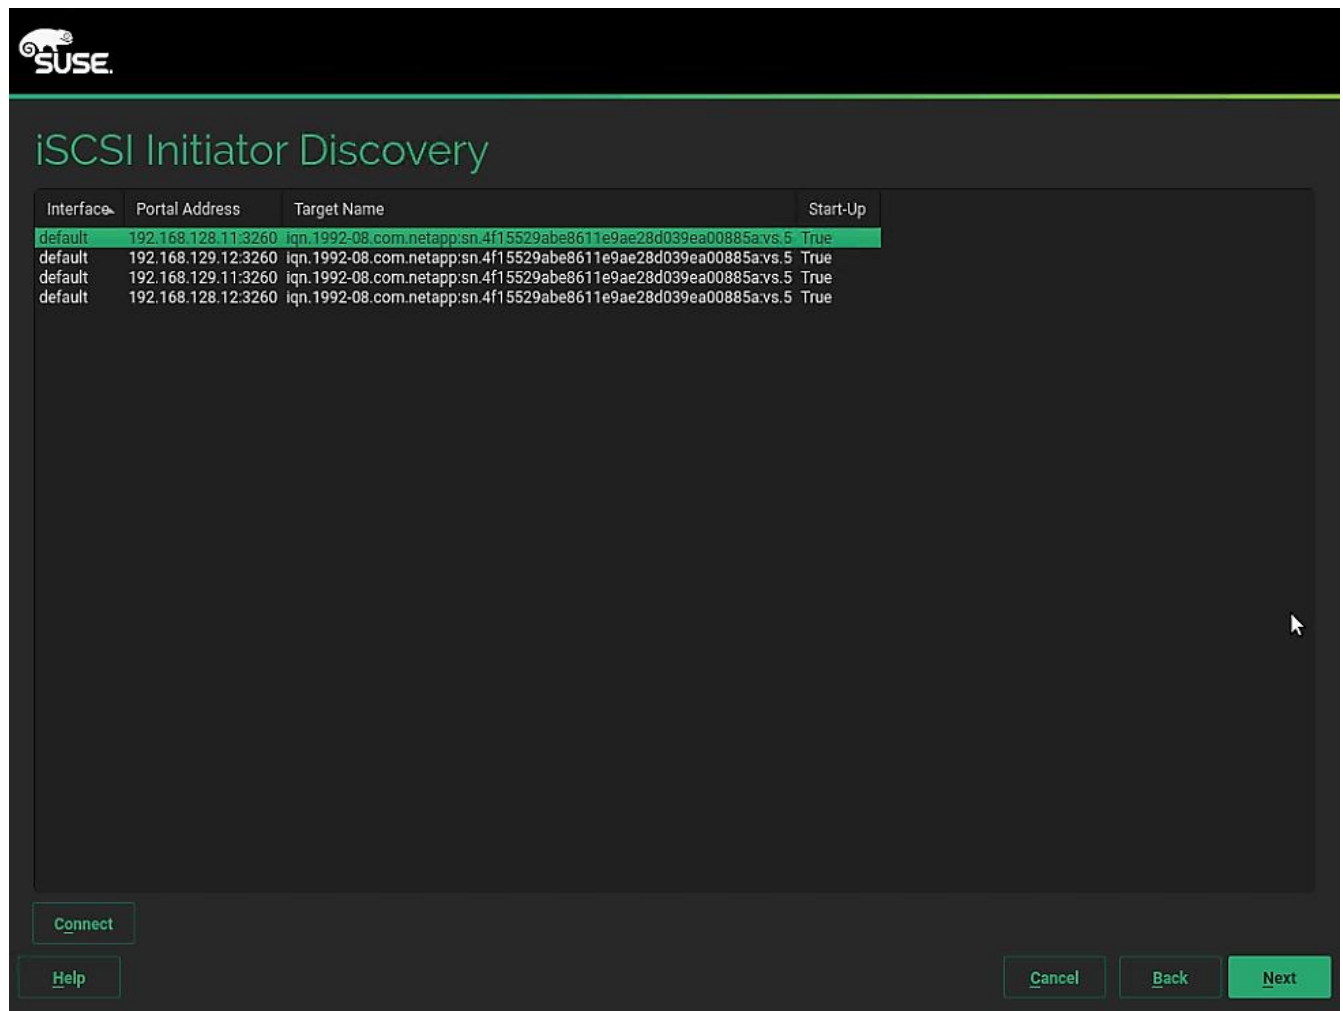
The screenshot shows the SUSE iSCSI Initiator Discovery window. At the top left is the SUSE logo. The title is "iSCSI Initiator Discovery". Below the title is a table with the following data:

Interface	Portal Address	Target Name	Start-Up
default	192.168.128.11:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	True
default	192.168.129.12:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	True
default	192.168.129.11:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	False
default	192.168.128.12:3260	iqn.1992-08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a.vs.5	True

A warning dialog box is displayed in the center of the window. The dialog has a title bar "Warning" and contains the text: "The target with this TargetName is already connected. Make sure that multipathing is enabled to prevent data corruption." Below the text are two buttons: "Continue" and "Cancel".

At the bottom of the window, there are several buttons: "Connect", "Help", "Cancel", "Back", and "Next".

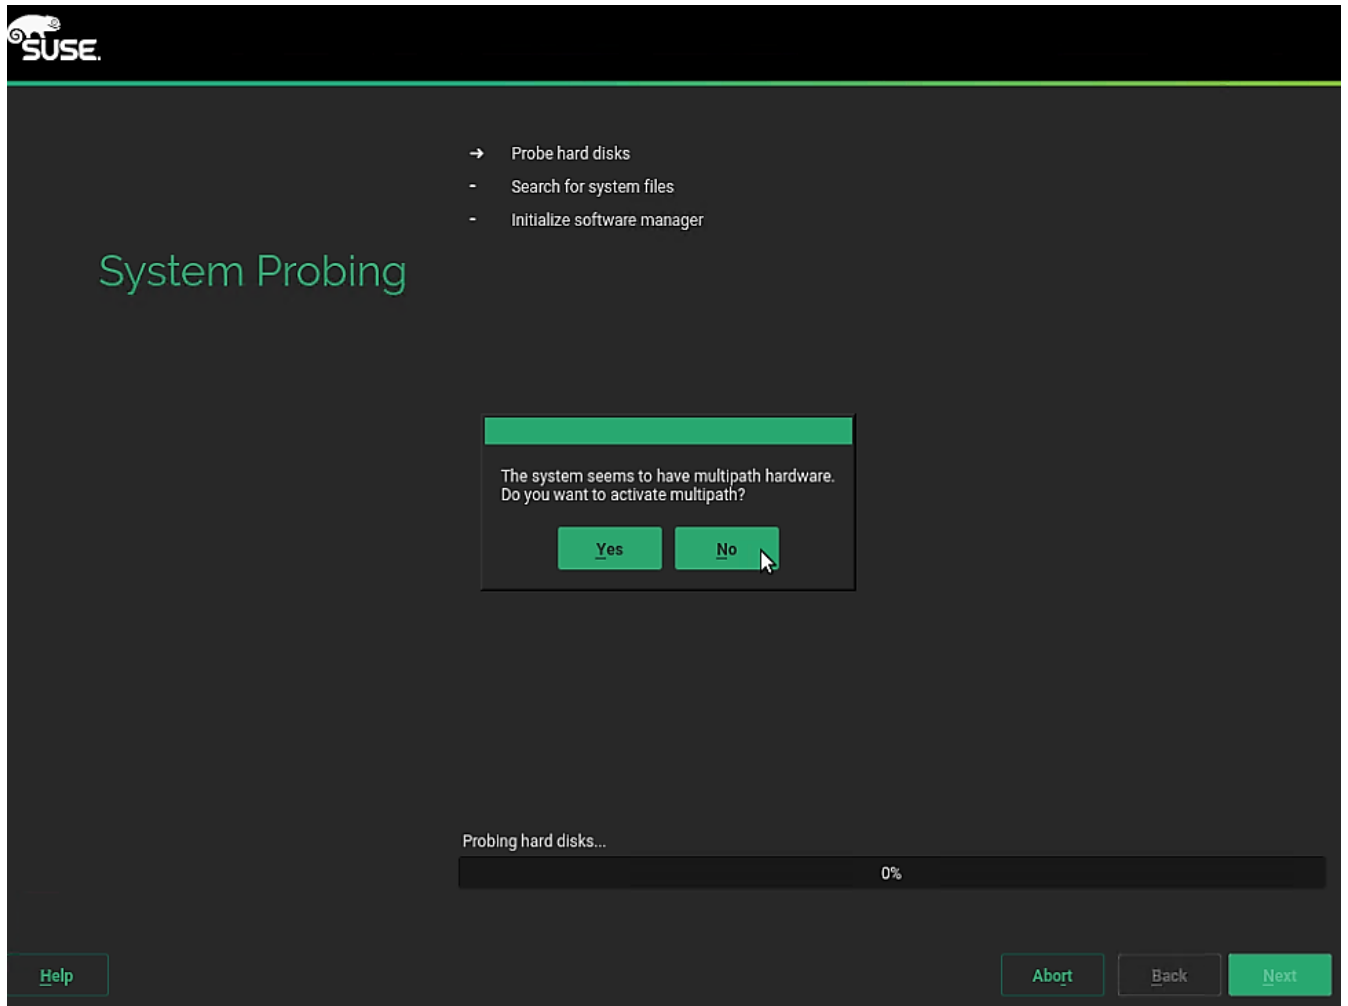
16. Have all paths discovered and startup set to value True.



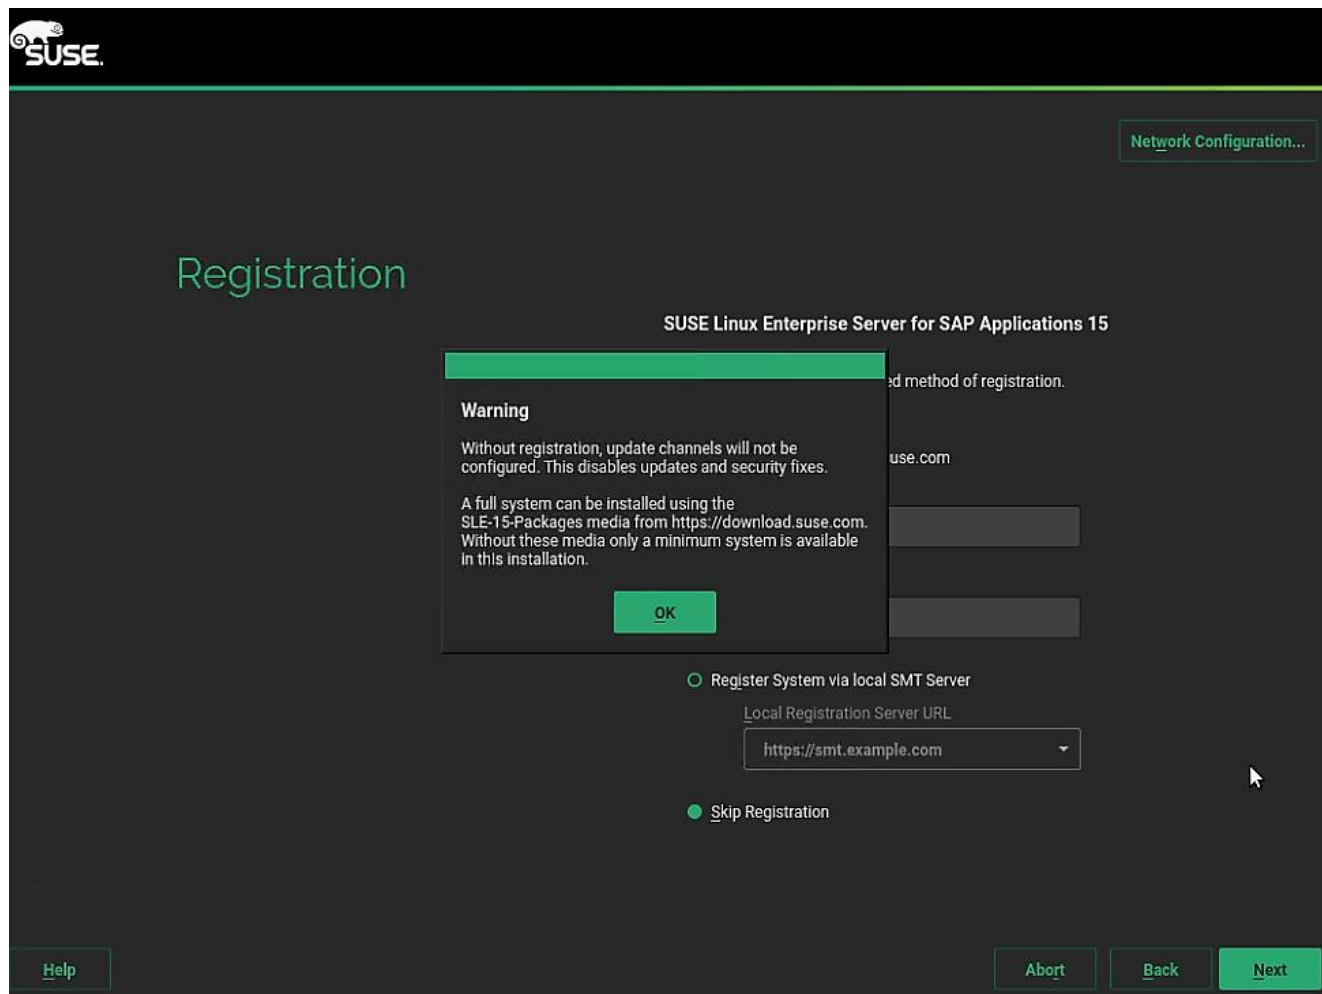
17. Click Next

18. Click Next on Disk Activation page.

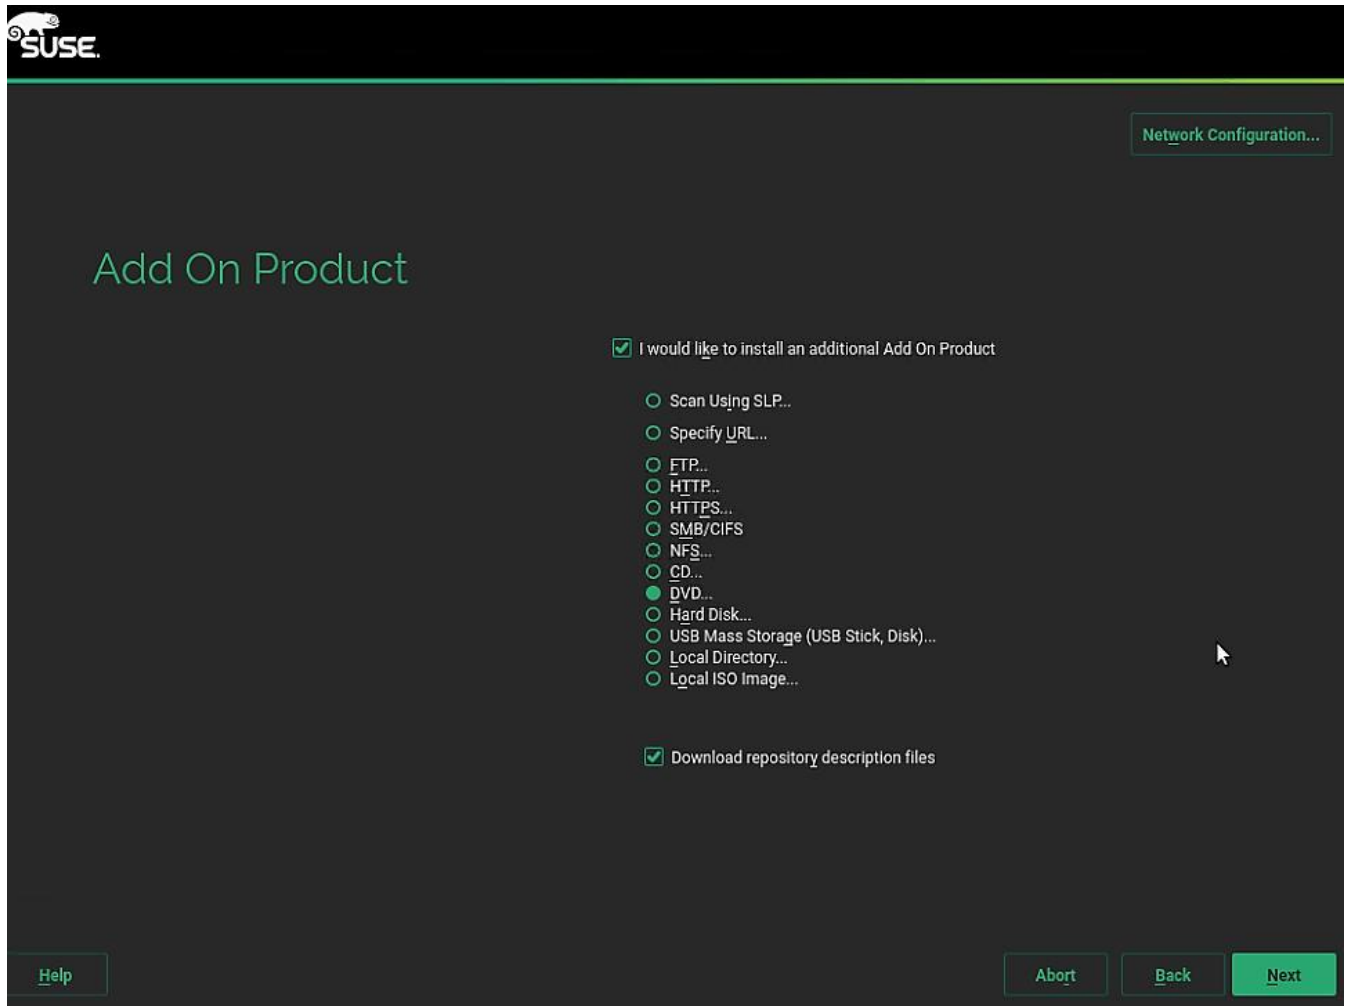
19. *System Probing* – Select No for the prompt to activate multipath. We will activate the multipath post installation.



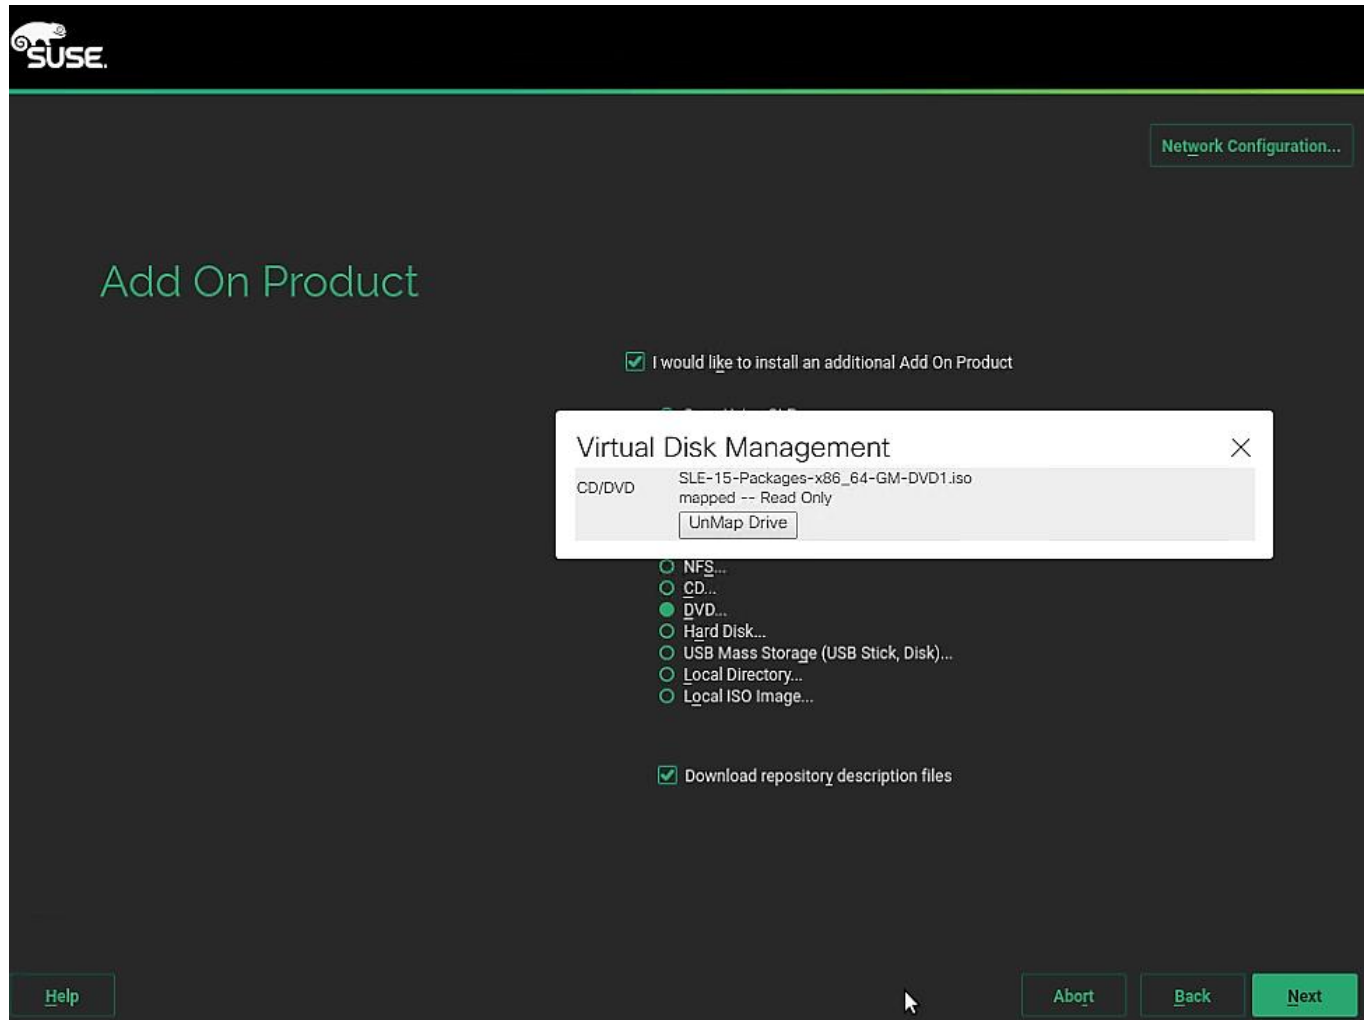
20. *Registration* – Select Skip Registration. We will register the system post installation. Click OK for the Warning.



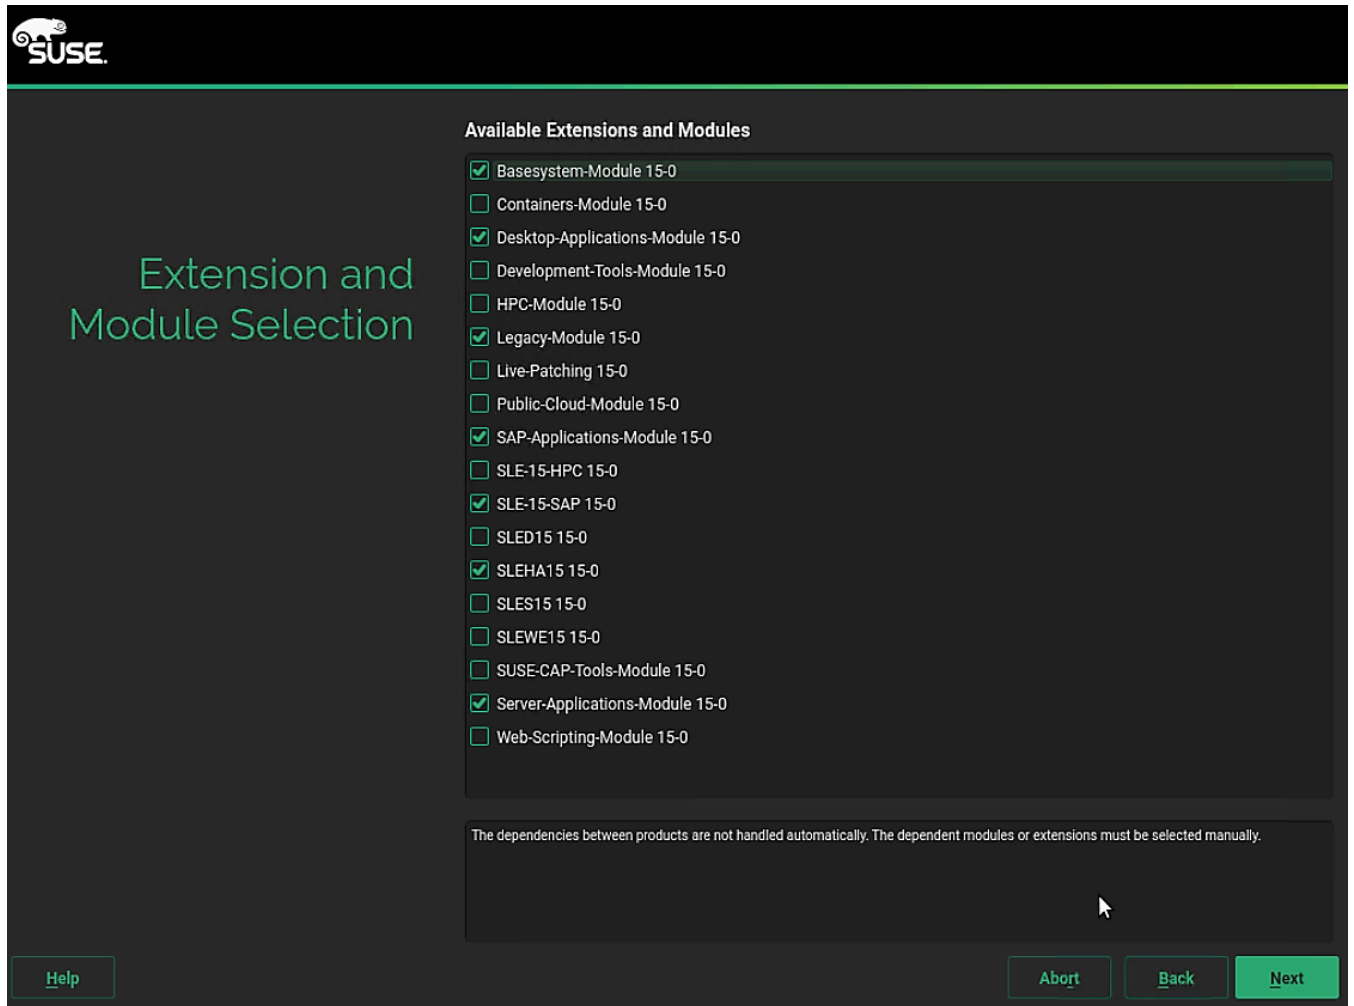
21. *Add On Product*: Select "I would like to install an additional Add On Product. And select DVD as the source. lick Next.



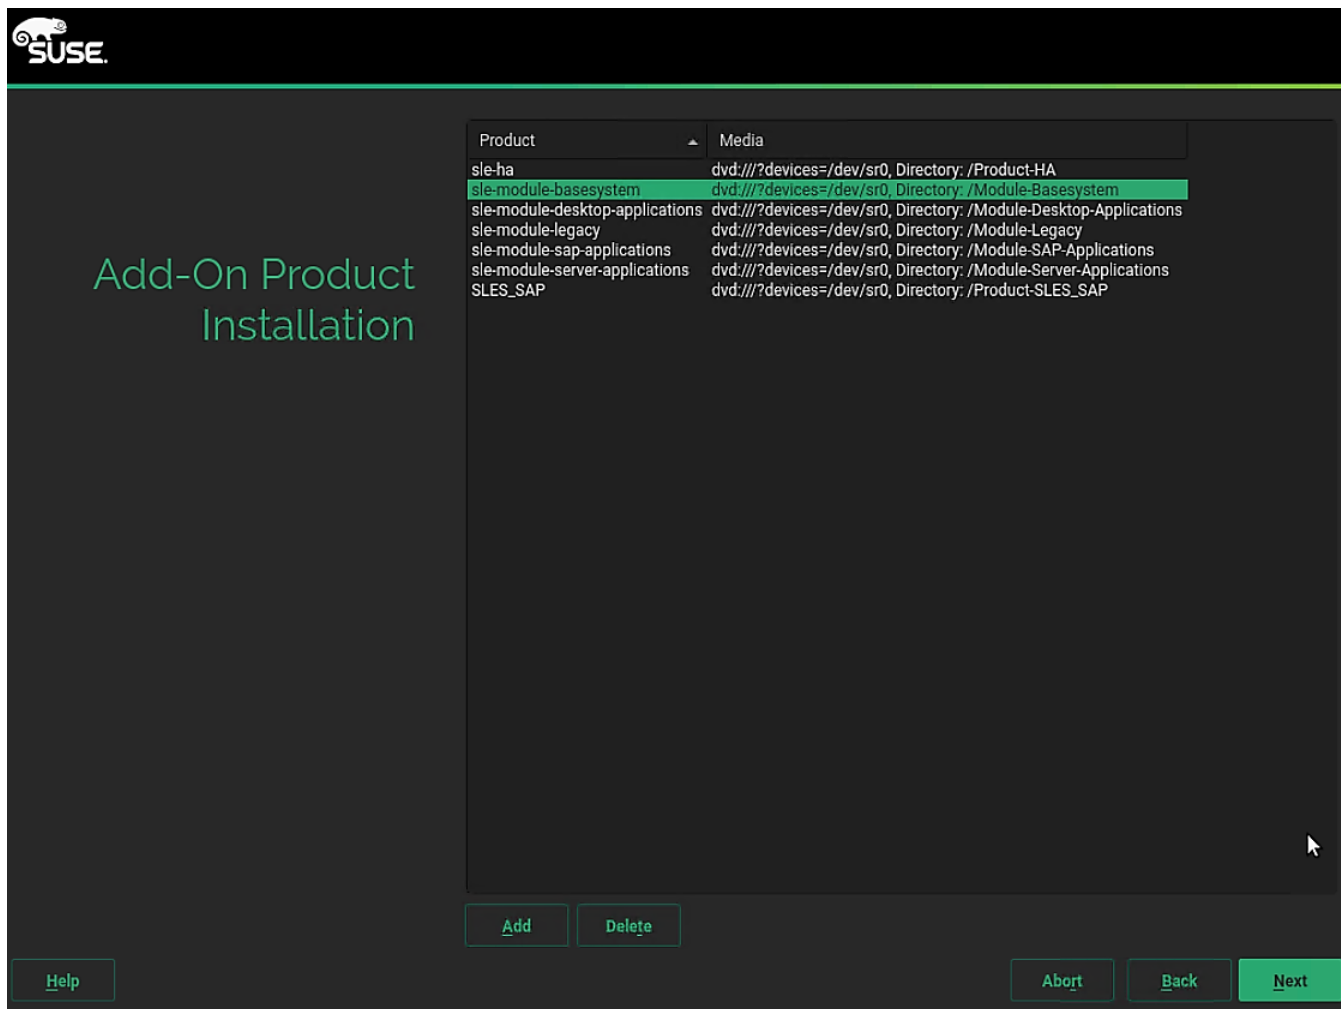
22. At this stage, unmap the currently mapped Installer DVD1 drive. Map the SLE-15-Packages-x86_64-GM-DVD1.iso Click Next.



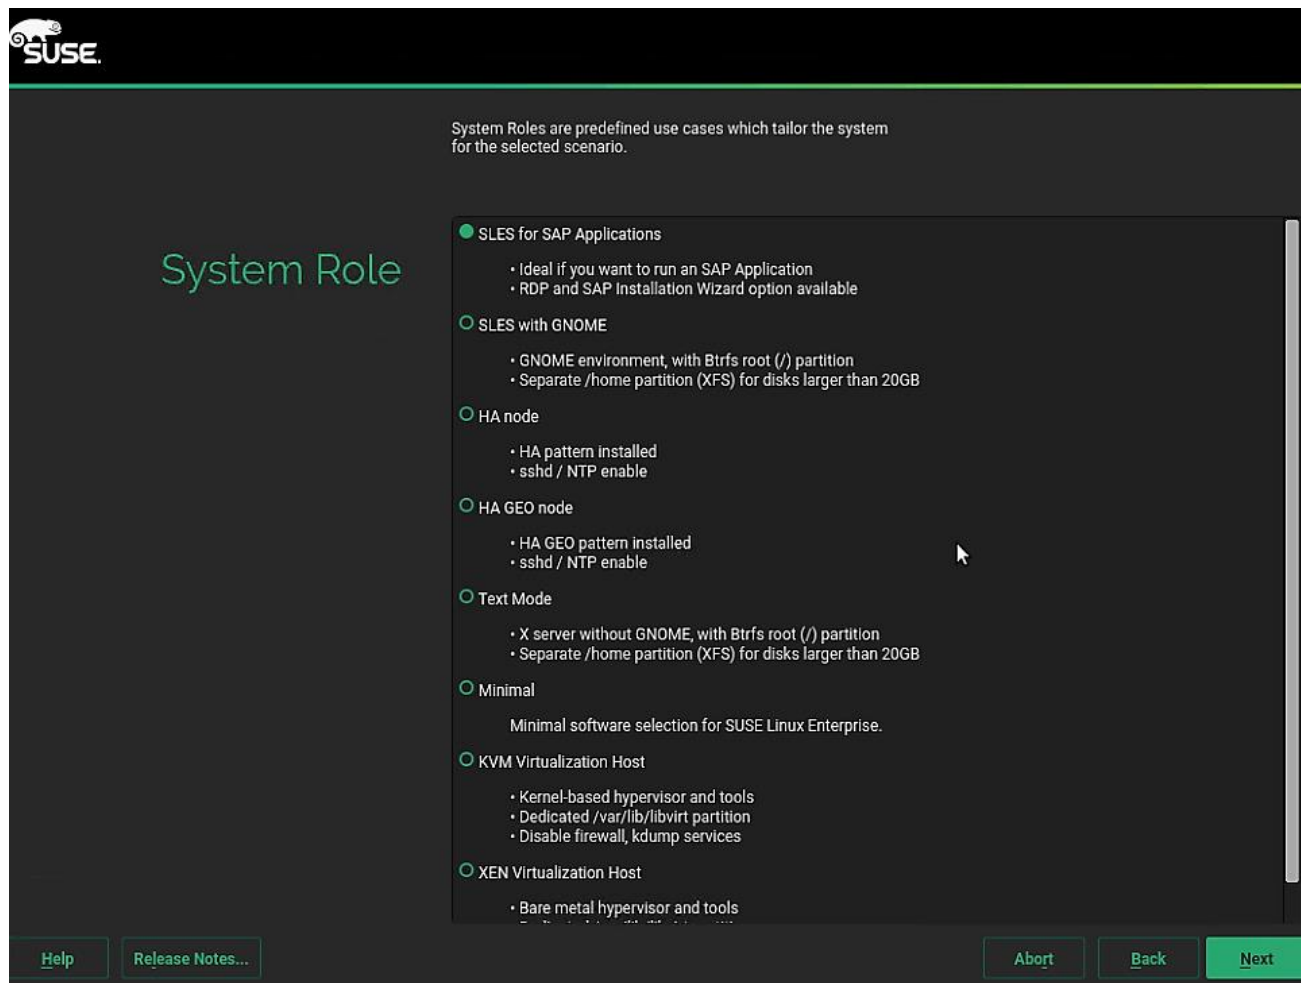
23. Extension and Module Selection – Select as appropriate. Click Next.



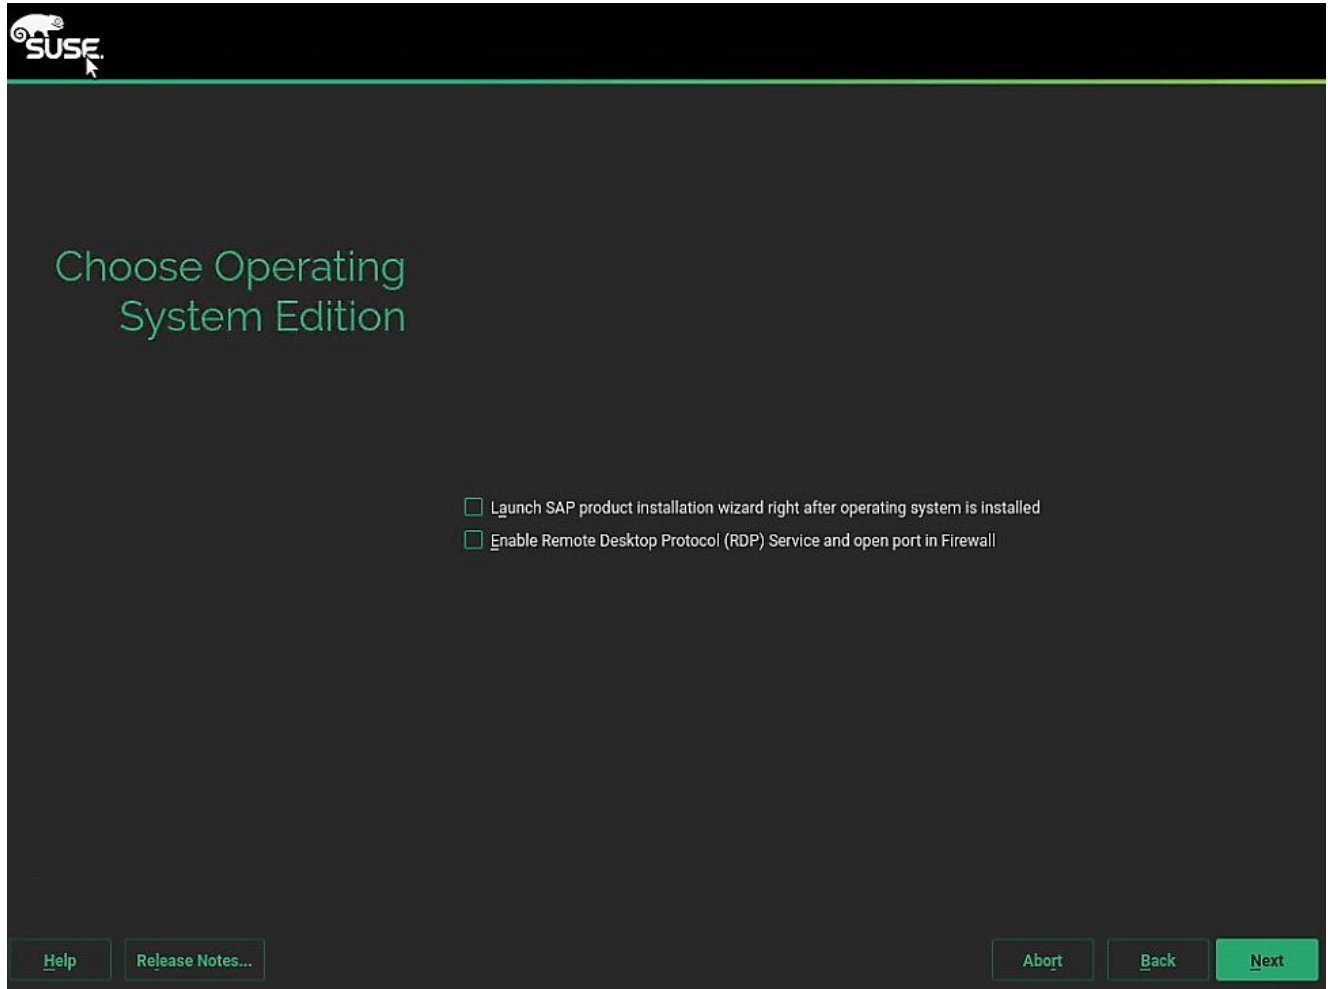
24. Click Next on the Add-On Product Installation page.



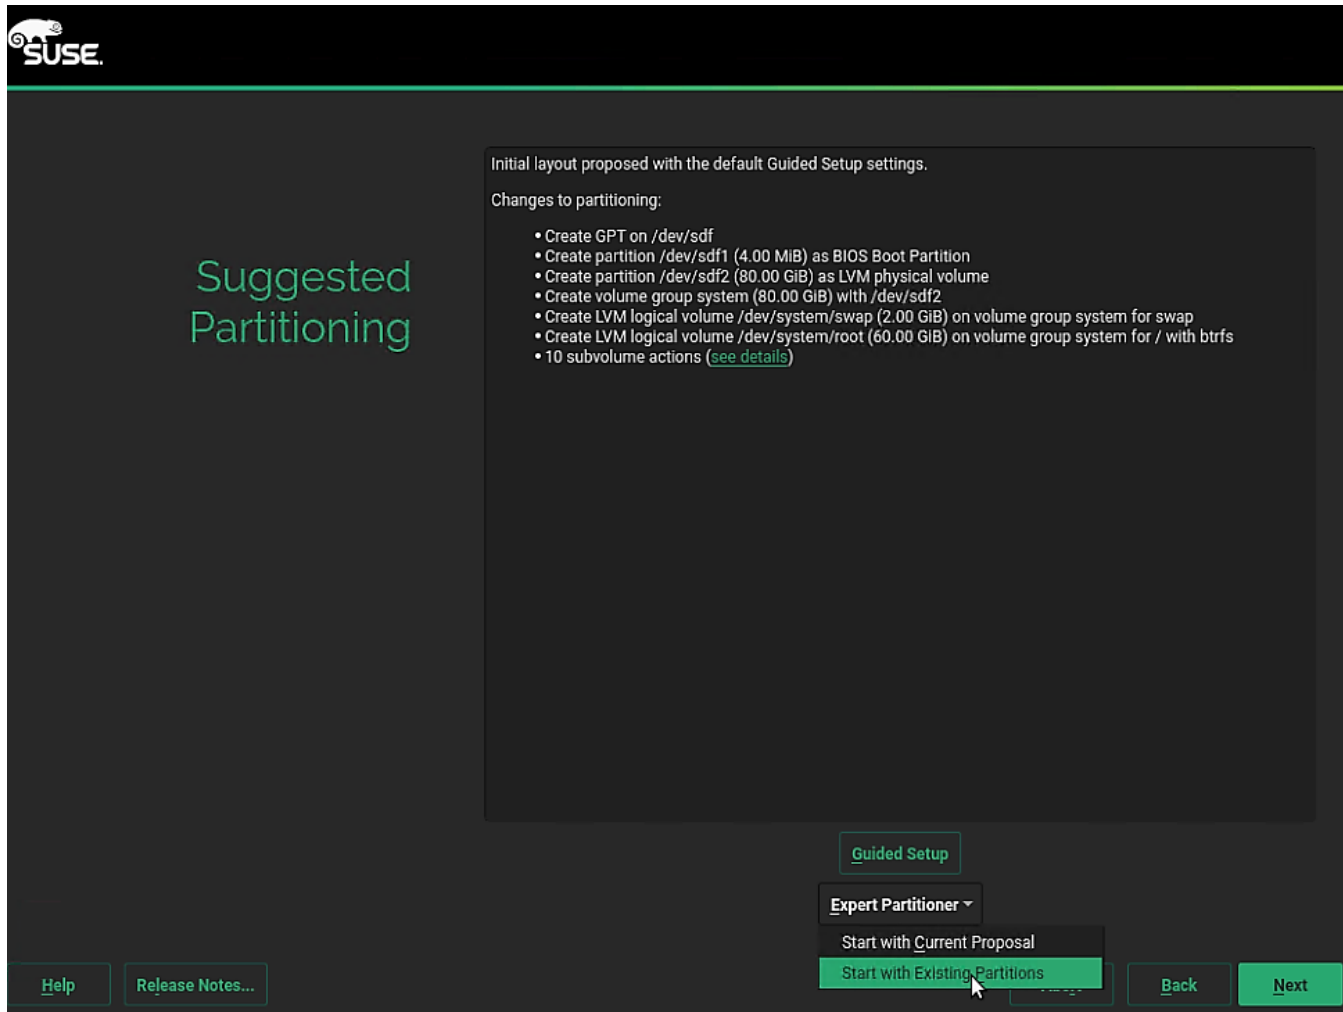
25. *System Role* - Keep the default selection of SLES for SAP Applications.



26. Choose Operating System Edition – Unselect both options. Click Next.

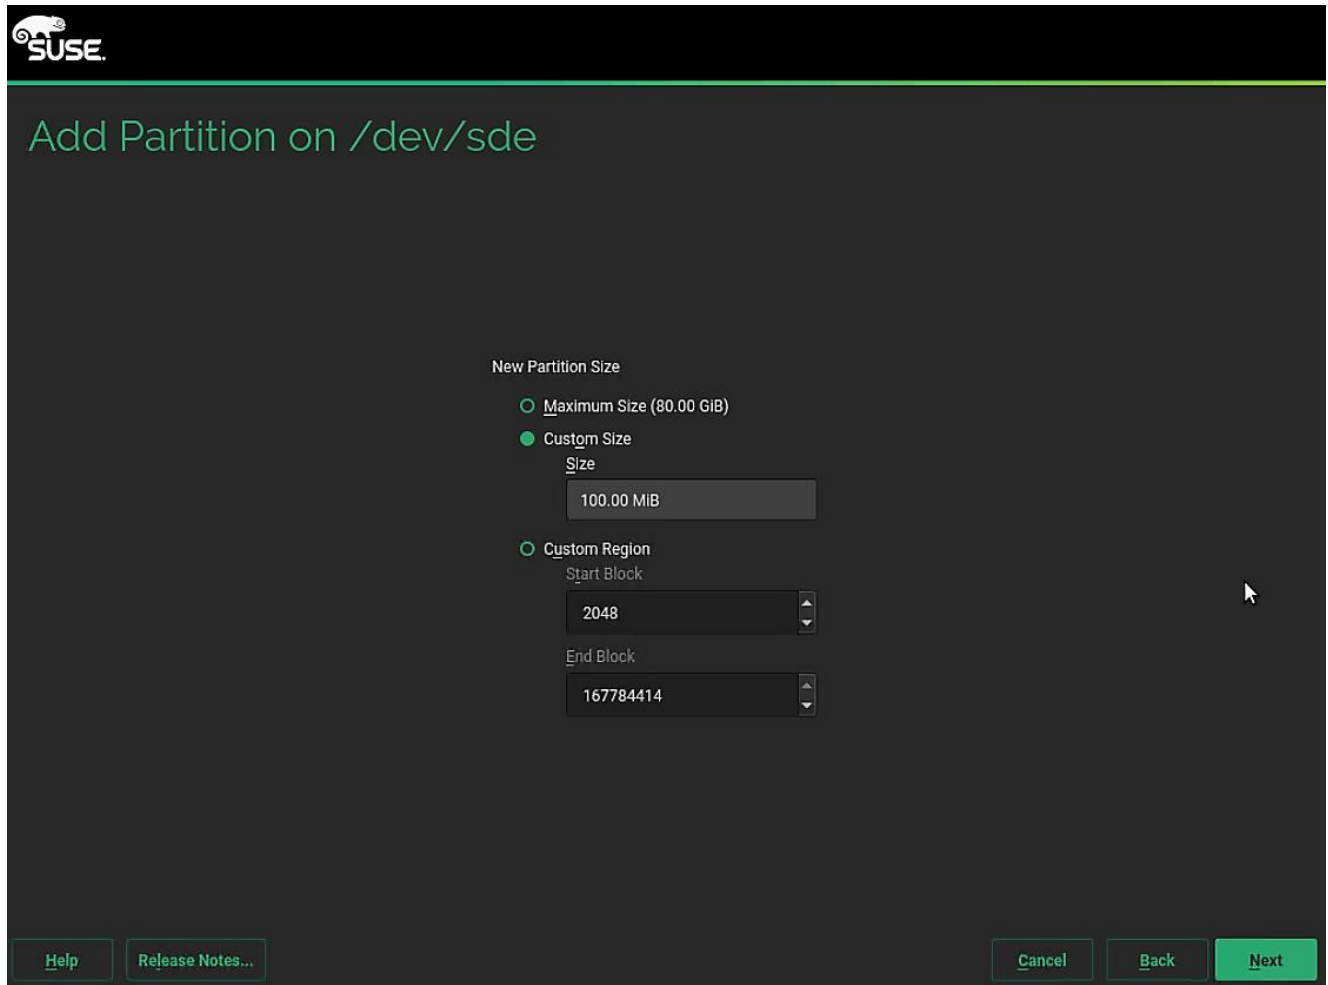


27. Suggested Partitioning – Click Expert Partitioner and Select Start with Existing Partitions.

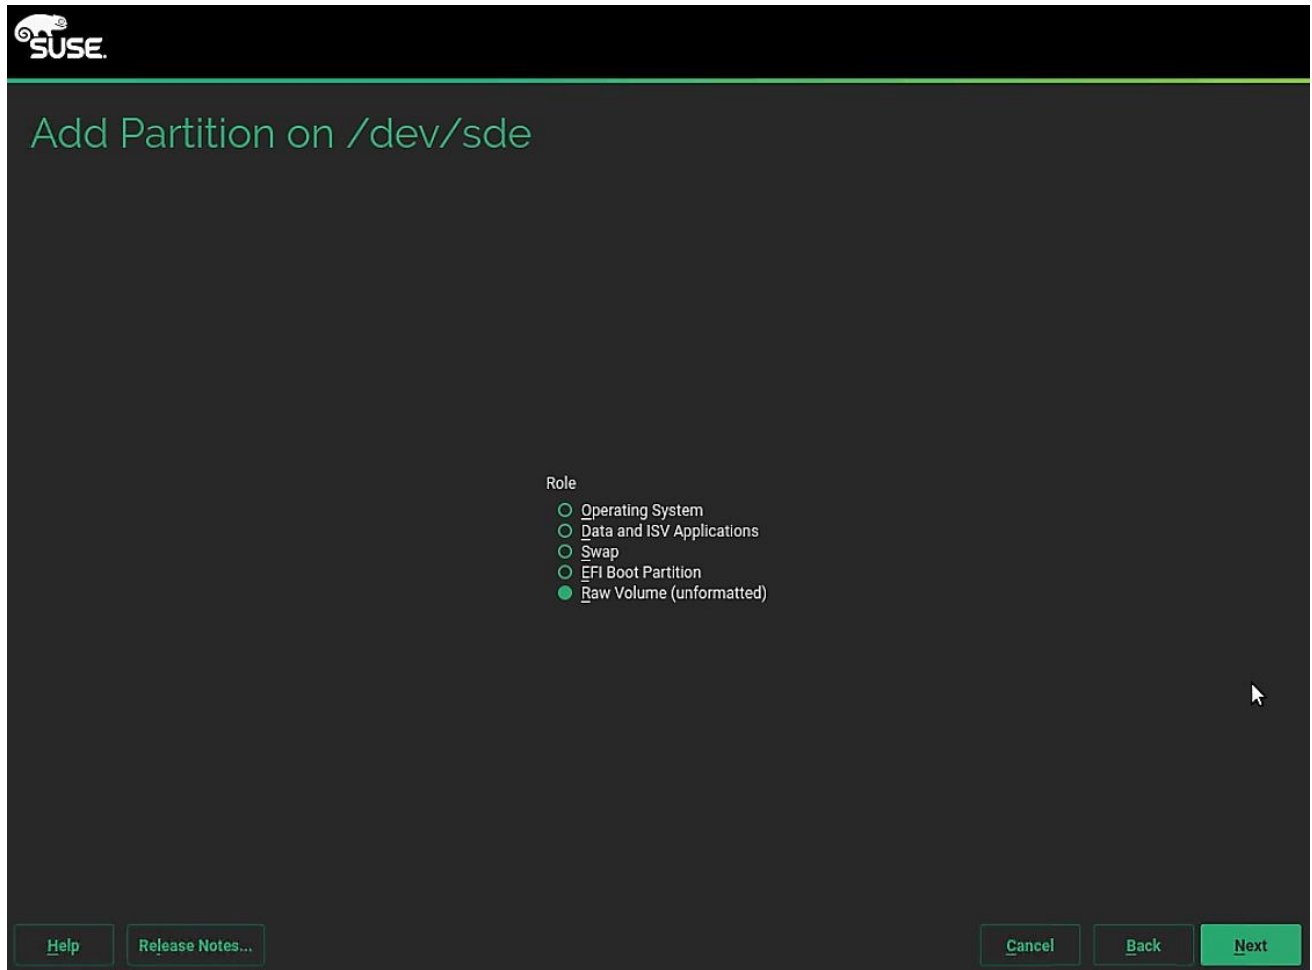


28. Expert Partitioner – Select one out of the four iSCSI NetApp devices listed. Click Add Partition.

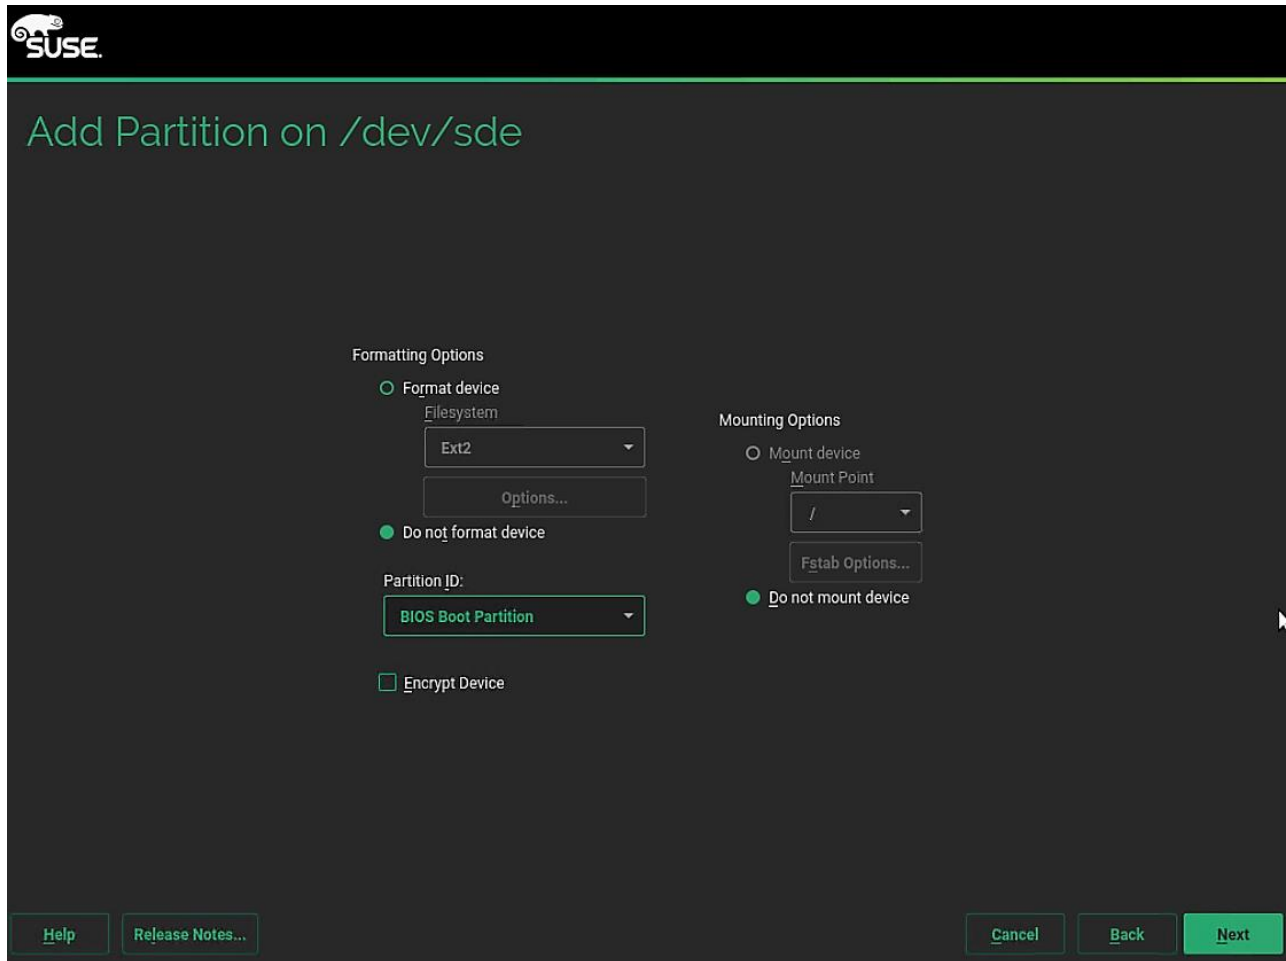
29. Create a 100 MiB size BIOS_Boot volume.



30. Select Raw Volume (unformatted) for the Role.



31. Select BIOS Boot Partition under Partition ID and leave the rest with default selections.



32. Add New Partition with Maximum Size, using up the rest of the space available on drive. Click Next.

Add Partition on /dev/sde

New Partition Size

Maximum Size (79.90 GiB)

Custom Size

Size

79.90 GiB

Custom Region

Start Block

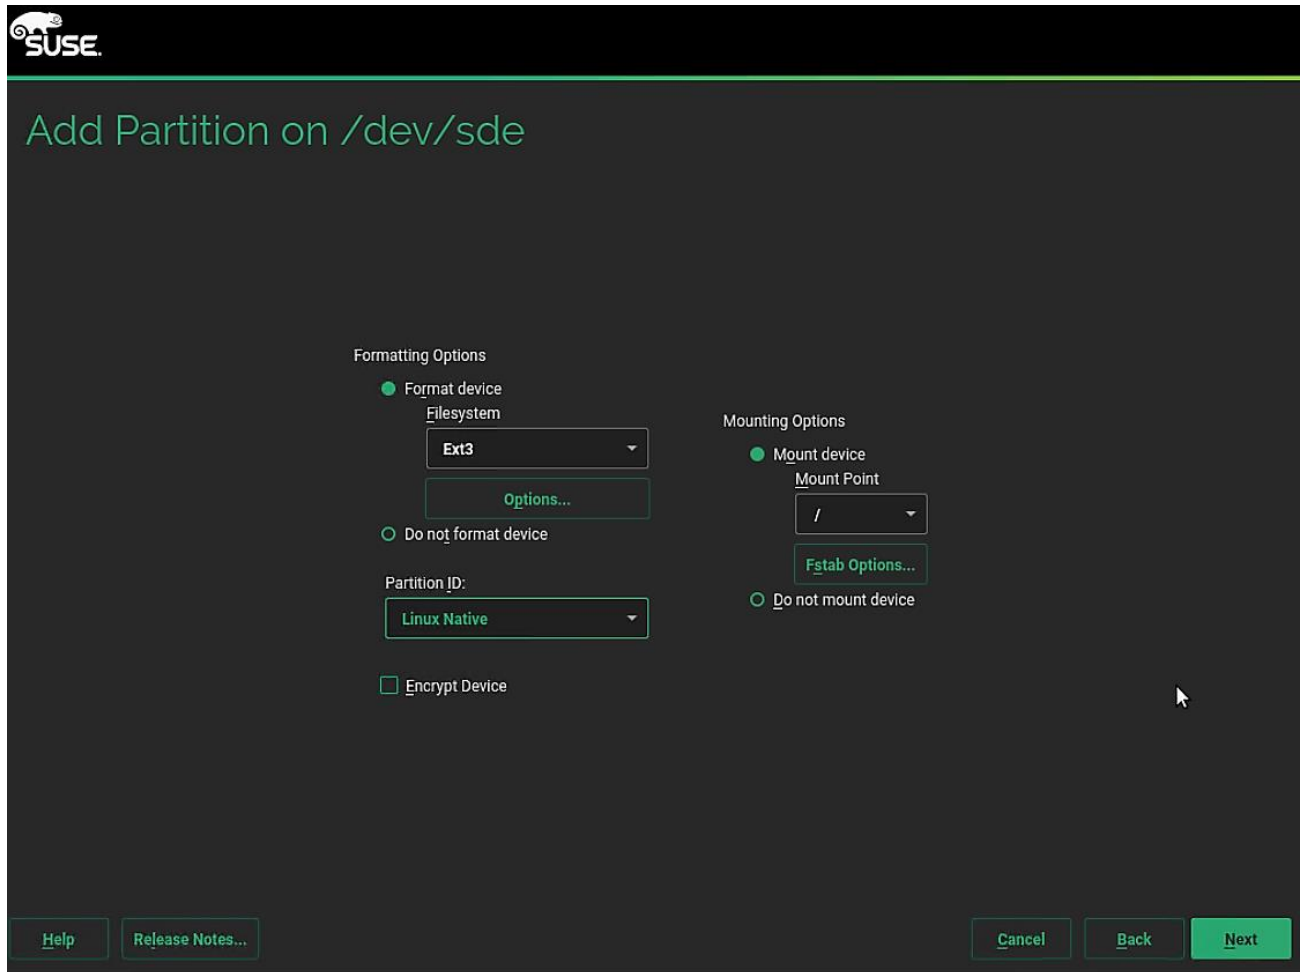
206848

End Block

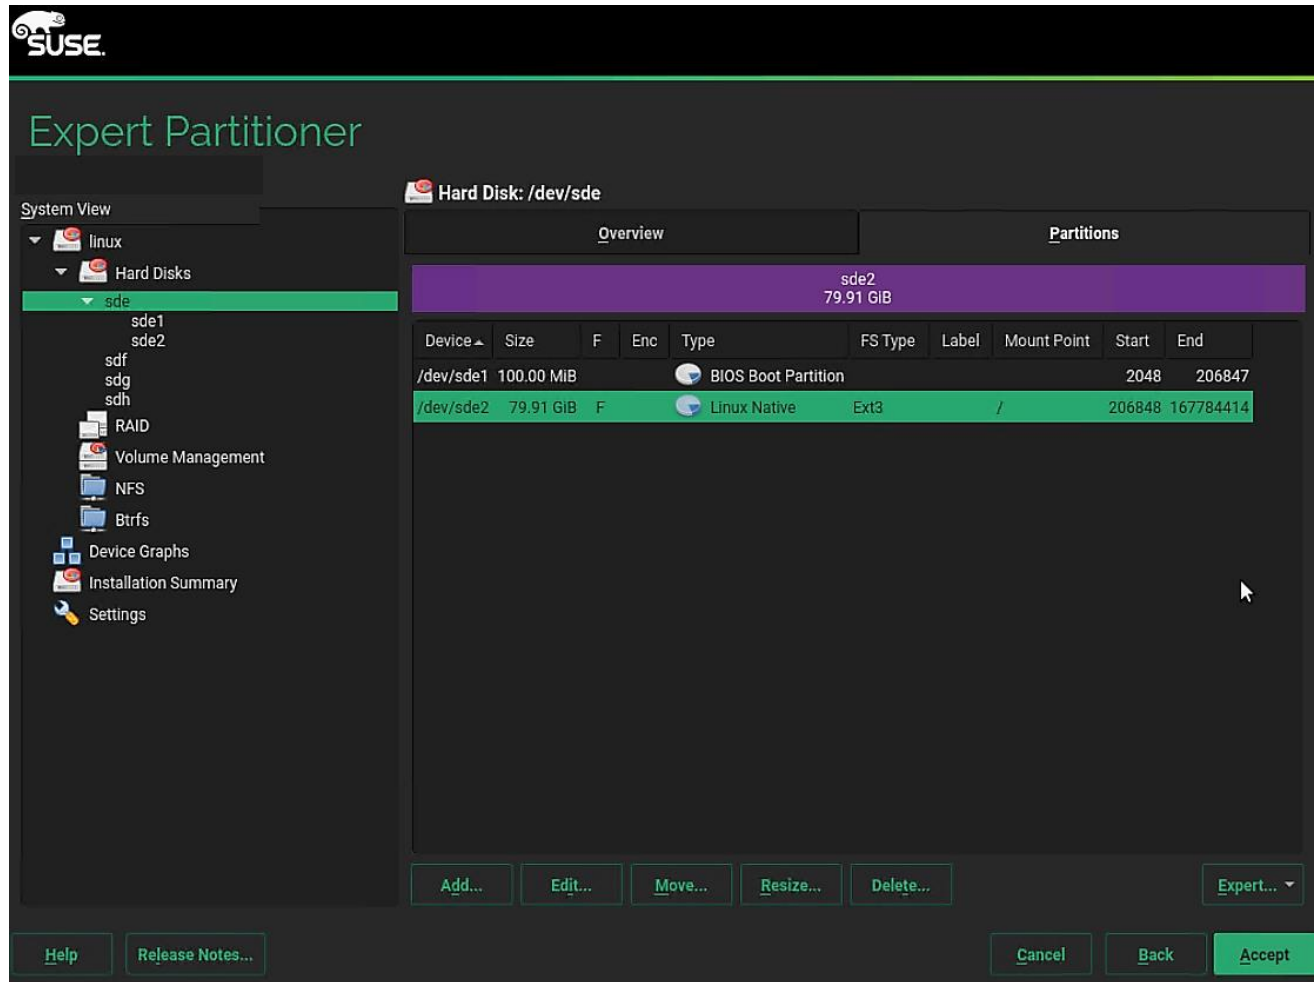
167784414

[Help](#) [Release Notes...](#) [Cancel](#) [Back](#) [Next](#)

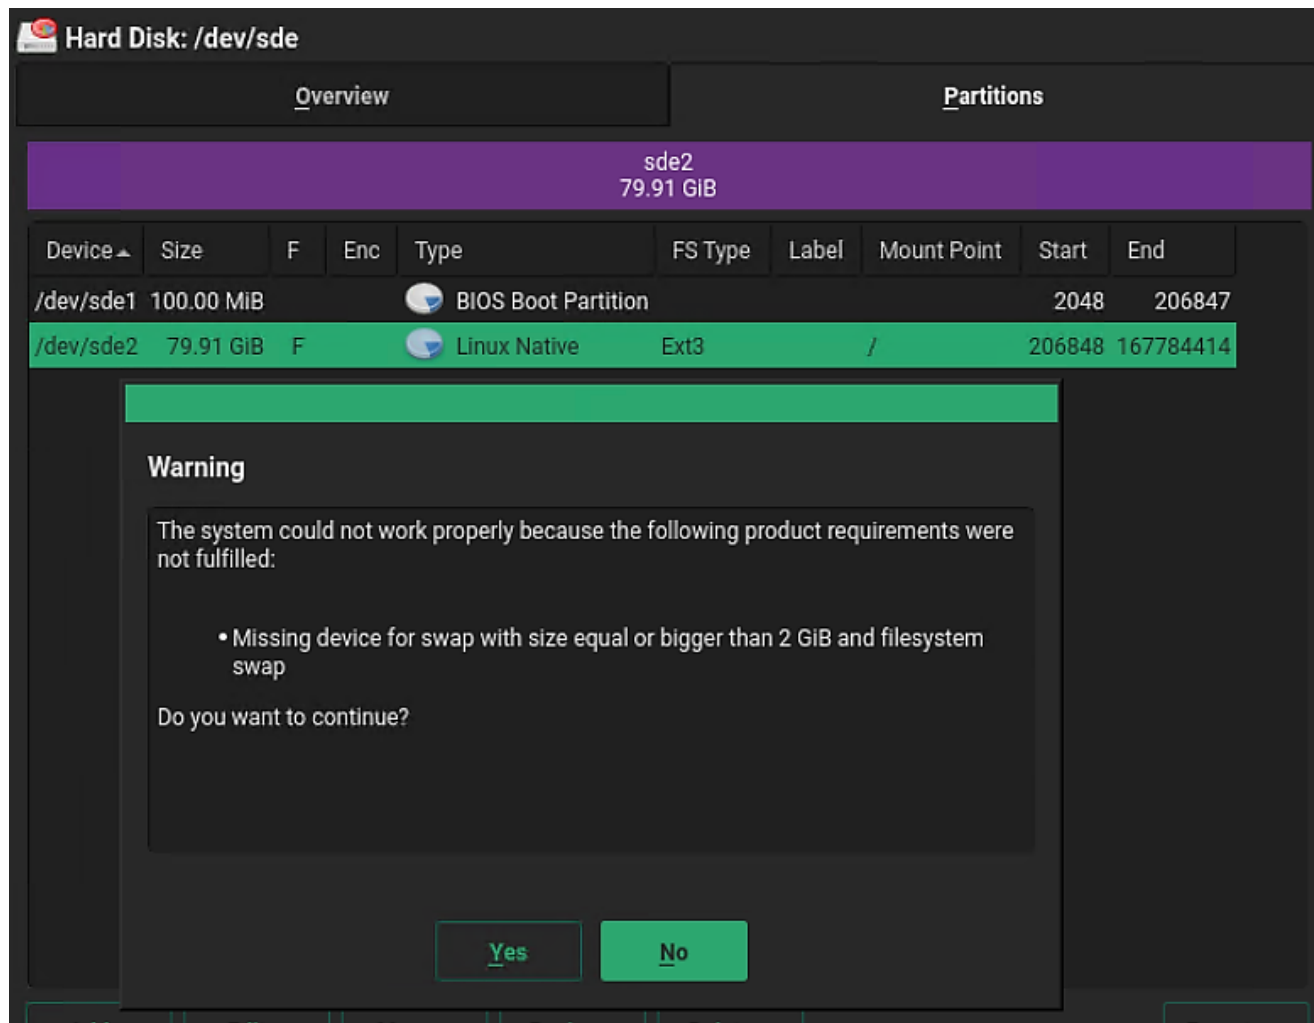
33. Select Operating System Role and click Next.



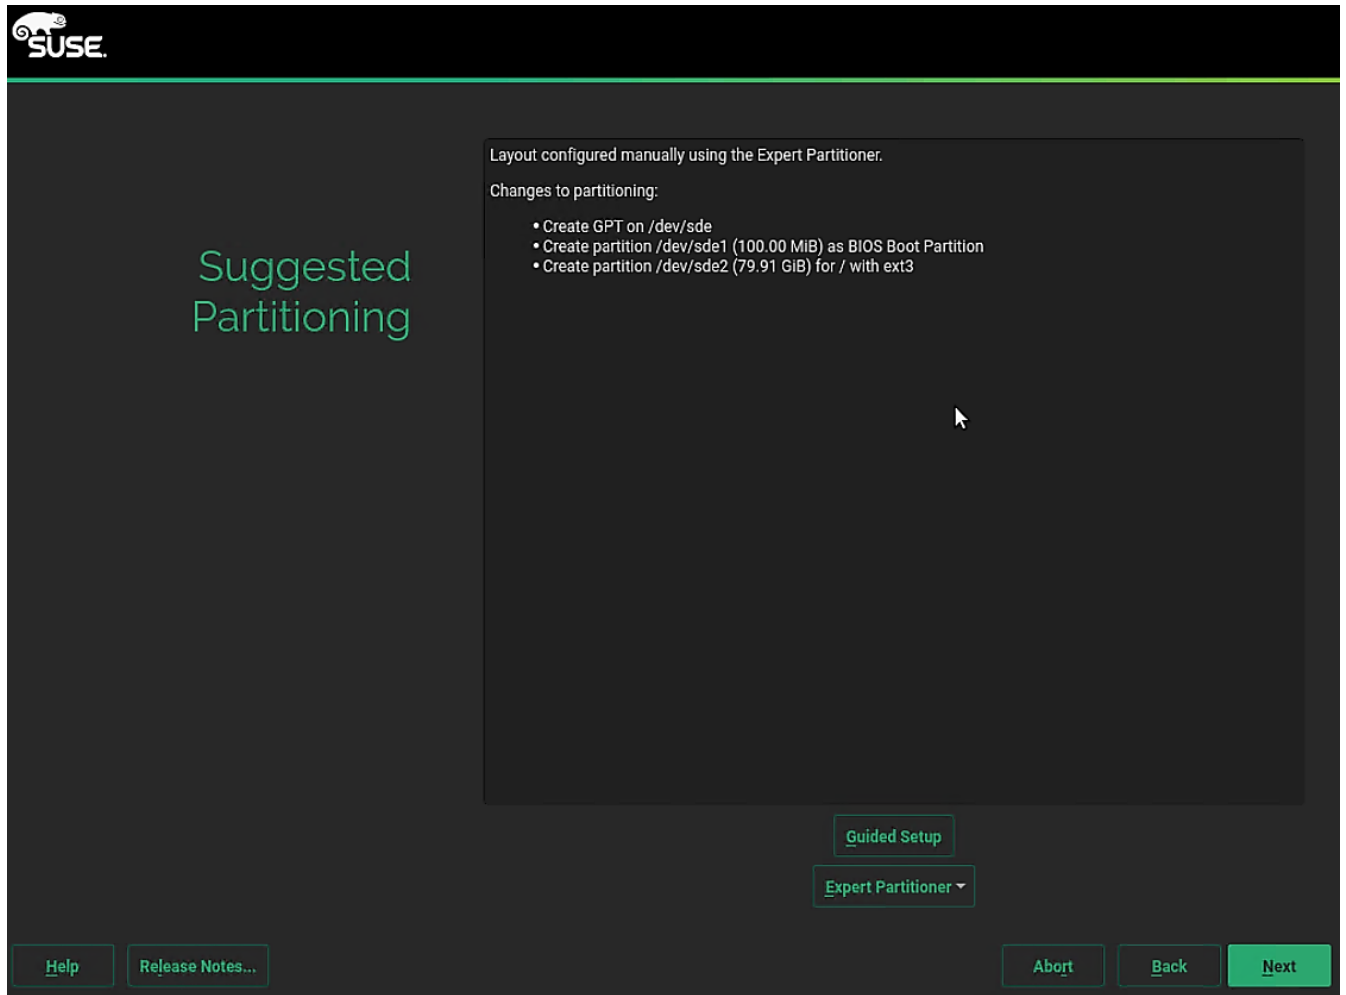
34. Select Ext3 for Filesystem. Click Next.



35. Click Accept. Click Yes for the warning, to continue. We will create the swap file later.



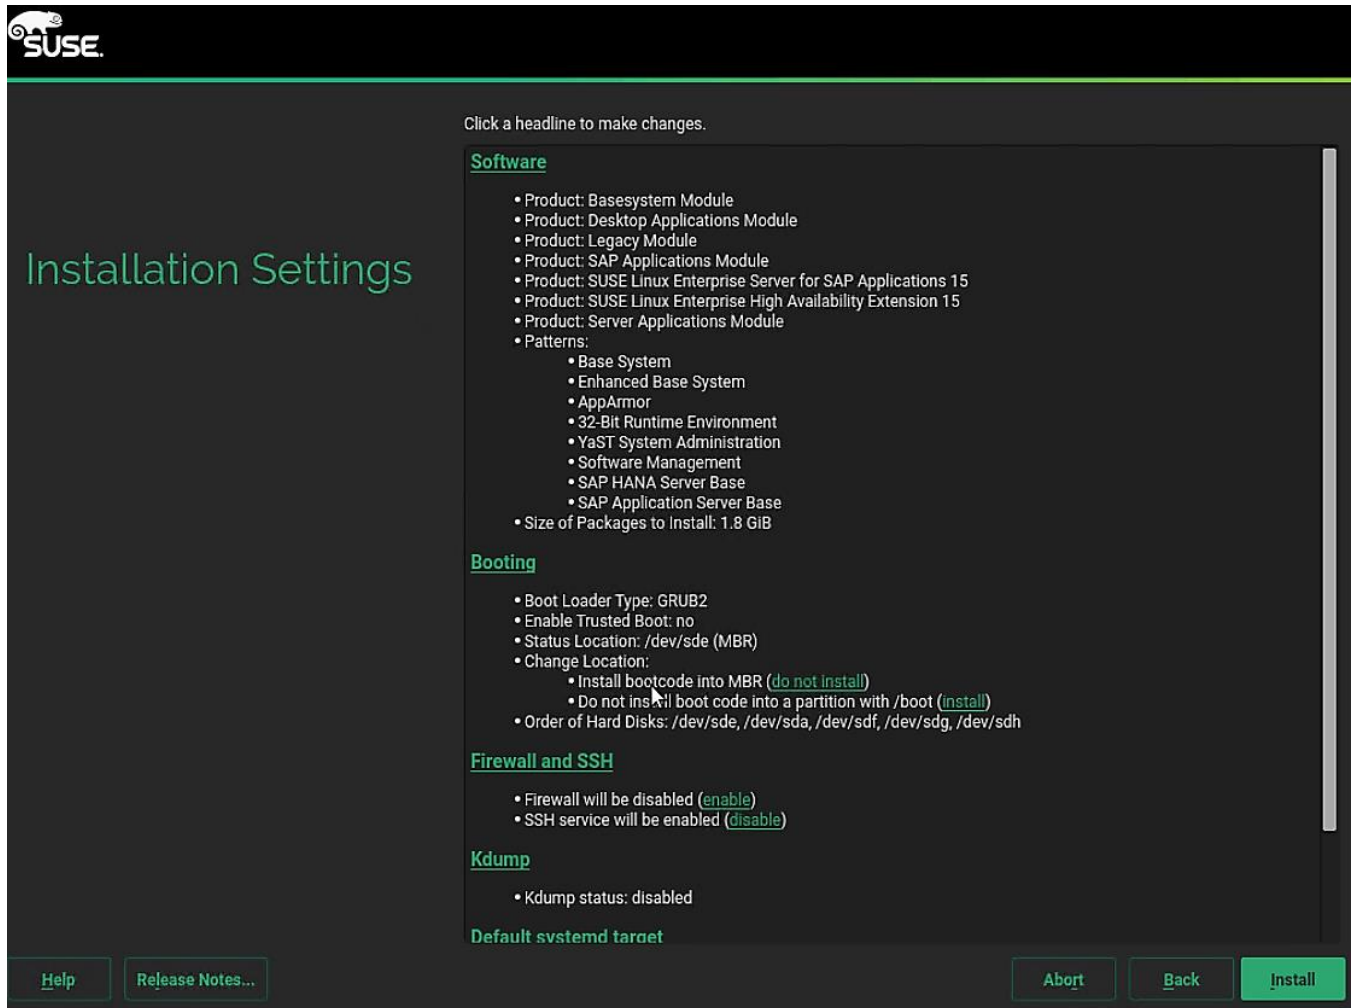
36. Suggested partitioning – Click Next.



37. Clock and Time Zone – choose the appropriate time zone and select Hardware clock set to UTC.

38. Password for the System Administrator “root” – Key in appropriate password <<var_sys_root-pw>>. Click Next.

39. On the Installation Settings screen, review the default information.



40. Click Install and select Install again for the subsequent 'Confirm Installation' prompt. The installation is started and you can monitor the status.

The screenshot shows the SUSE installation progress window. On the left, the text "Performing Installation" is displayed in green. The main area is divided into three tabs: "Details", "SLEHA15 Release Notes", and "SLE-15-SAP Release Notes". The "Details" tab is active and contains a table with the following data:

Media	Remaining	Packages	Time
Total	1.772 GiB	1042	
SLES15-15-0			
Medium 1	637.6 KiB		
sle-module-basesystem			
Medium 1	1.681 GiB	949	
sle-module-desktop-applications			
Medium 1	54.35 MiB		
sle-module-legacy			
Medium 1	2.4 KiB		
sle-module-sap-applications			
Medium 1	23.74 MiB		
SLES SAP			

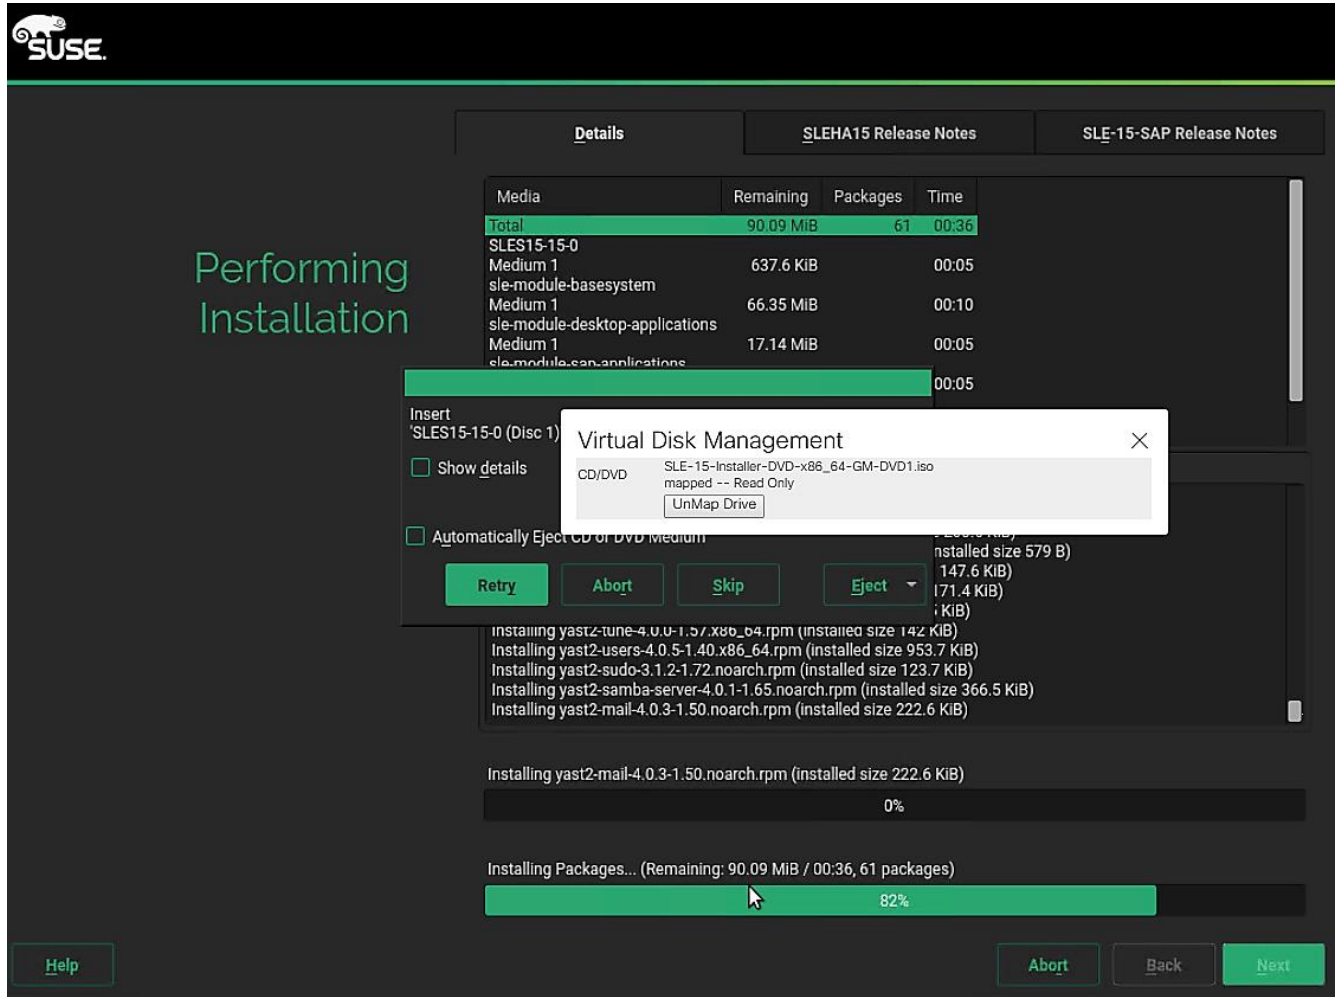
Below the table, the "Actions performed:" section lists the following actions:

- Installing qt59-data-3.20.30-2.13.noarch.rpm (installed size 3.8 KiB)
- Installing fire3-subvna-3.22.30-2.11.noarch.rpm (installed size 8.2 KiB)
- Installing qt59-translation-theme-0.17.1.21.noarch.rpm (installed size 76.7 KiB)
- Installing libd-legacy-2.0.4.6-1.1.noarch.rpm (installed size 6.7 KiB)
- Installing libX11-data-1.6.5-1.41.noarch.rpm (installed size 1.22 MiB)
- Installing libicu60_2-ledata-60.2-1.17.noarch.rpm (installed size 25.65 MiB)
- Installing libnl-config-3.3.0-1.29.noarch.rpm (installed size 2.6 KiB)
- Installing libstorage-ng-lang-3.3.300-1.1.noarch.rpm (installed size 679.5 KiB)
- Installing libthai-data-0.1.27-1.16.x86_64.rpm (installed size 558.6 KiB)
- Installing libtirpc-netconfig-1.0.2-1.127.x86_64.rpm (installed size 2 KiB)
- Installing patterns-base-32bit-20171206-22.1.x86_64.rpm (installed size 51 B)
- Installing patterns-base-apparmor-32bit-20171206-22.1.x86_64.rpm (installed size 54 B)

At the bottom of the window, there are three progress bars and buttons:

- A green progress bar showing 100% completion for "Installing patterns-base-apparmor-32bit-20171206-22.1.x86_64.rpm (installed size 54 B)".
- A green progress bar showing 16% completion for "Installing Packages... (Remaining: 1.772 GiB, 1042 packages)".
- Buttons for "Help", "Abort", "Back", and "Next".

41. When prompted during the installation process, remap installation DVD1.



42. After the installation is complete, a reboot alert appears. The system will reboot and boot from disk on startup. Login using the root.

```

Welcome to SUSE Linux Enterprise Server for SAP Applications 15 (x86_64) - Kernel 4.12.14-23-default (tty1).

eth0:
eth1:
eth3:
eth4:
eth5:
eth7:

linux-ngcg login: root
Password:
Last login: Thu Oct 10 14:22:42 on tty1
linux-ngcg:~ #
    
```

Post Installation Steps

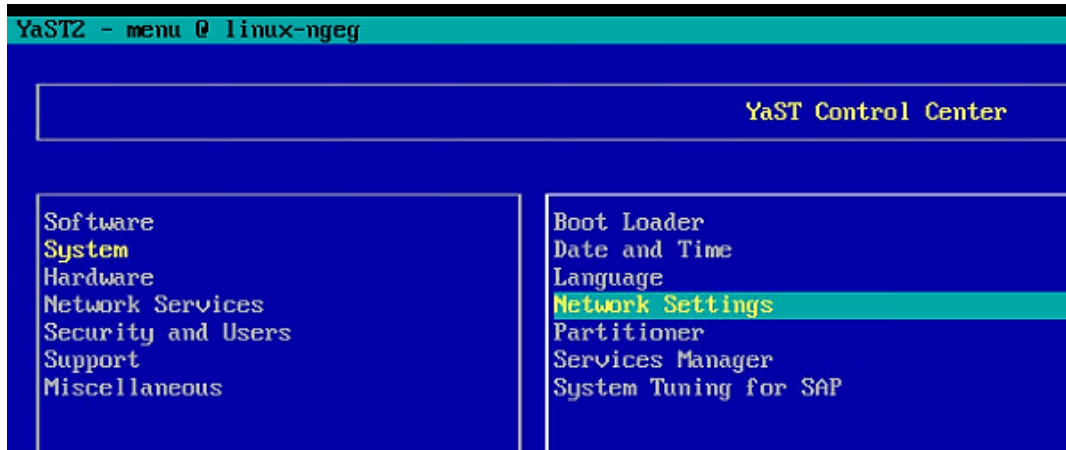
Configure the Network

As part of the post-install configuration, you will define the hostname, network configuration, kernel/packages update, and tuning as per the SAP Notes recommendations.

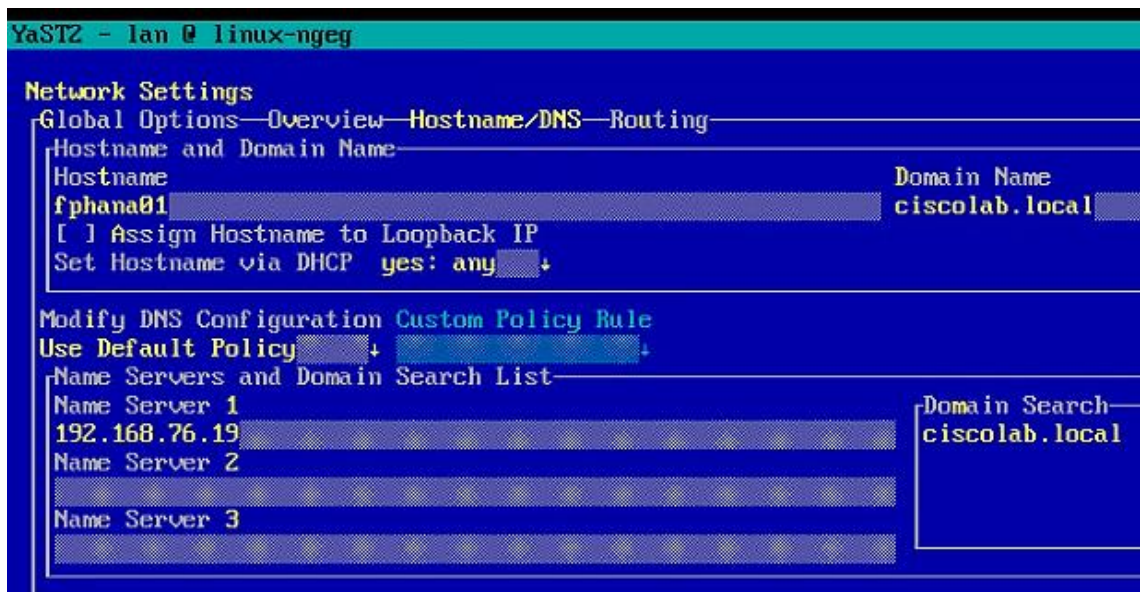
To configure the network, follow these steps:

1. Configure the hostname and disable IPV6.

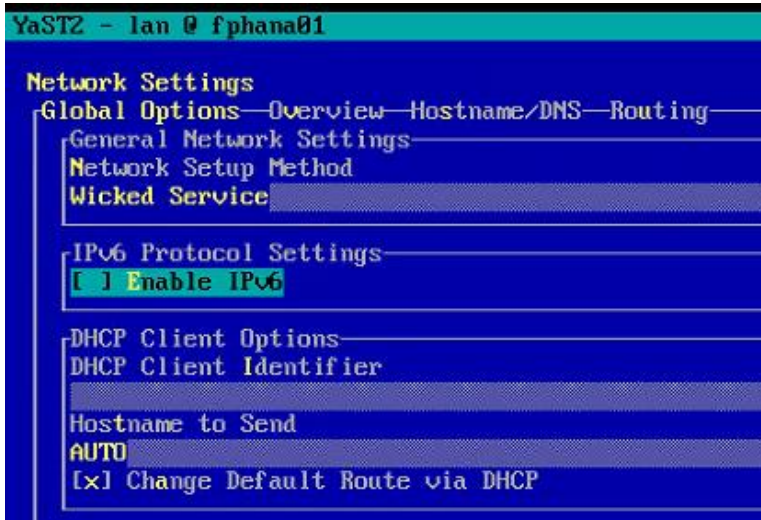
```
#yast2
```



2. System > Network Settings and select Run > Alt+s to select Hostname/DNS tab.



3. Input the <<var_hostname.domain>>. Also, key in DNS server address of your network for resolution, if any and select Alt+o.
4. On the Global Options tab with Alt+g, you can choose to disable IPV6, by unselecting the Enable IPV6 option as shown in the figure below. Changing the IPV6 setting requires a reboot to effect the changes.



5. Select Alt+o to save the Network Configuration. Select Alt+q to quit the YaST Control center.
6. Perform a reboot to effect the IPV6 selection and also the hostname settings.

```
#reboot
```

7. Host networking configuration:
 - a. The vNIC to MAC address mapping information for a host can be obtained from the network tab of that host's Service Profile.

vNICs	
Name	MAC Address
vNIC HANA-AppServer	00:25:B5:0A:00:0B
vNIC HANA-data	00:25:B5:0A:00:08
vNIC HANA-internode	00:25:B5:0B:00:06
vNIC HANA-log	00:25:B5:0B:00:07
vNIC HANA-node-admin	00:25:B5:0A:00:09
vNIC HANA-node-backup	00:25:B5:0B:00:09
vNIC iSCSI-A	00:25:B5:0A:00:0A
vNIC iSCSI-B	00:25:B5:0B:00:08

- b. At the host OS level, the Ethernet interface to MAC address mapping can be ascertained with the 'ip address' command.

```

Welcome to SUSE Linux Enterprise Server for SAP Applications 15 (x86_64) - Kernel 4.12.14-23-default (tty1).

eth0:
eth1:
eth3:
eth4:
eth5:
eth7:

fphana01 login: root
Password:
Last login: Fri Oct 11 07:11:03 on tty1
fphana01:~ # ip address
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    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> mtu 9000 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0b:00:06 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0a:00:09 brd ff:ff:ff:ff:ff:ff
4: ibft0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0a:00:0a brd ff:ff:ff:ff:ff:ff
    inet 192.168.128.52/24 brd 192.168.128.255 scope global ibft0
        valid_lft forever preferred_lft forever
5: eth3: <BROADCAST,MULTICAST> mtu 9000 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0a:00:0b brd ff:ff:ff:ff:ff:ff
6: eth4: <BROADCAST,MULTICAST> mtu 9000 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0a:00:08 brd ff:ff:ff:ff:ff:ff
7: eth5: <BROADCAST,MULTICAST> mtu 9000 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0b:00:07 brd ff:ff:ff:ff:ff:ff
8: ibft1: <BROADCAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP group default qlen 1000
    link/ether 00:25:b5:0b:00:08 brd ff:ff:ff:ff:ff:ff
    inet 192.168.129.52/24 brd 192.168.129.255 scope global ibft1
        valid_lft forever preferred_lft forever
9: eth7: <BROADCAST,MULTICAST> mtu 9000 qdisc noop state DOWN group default qlen 1000
    link/ether 00:25:b5:0b:00:09 brd ff:ff:ff:ff:ff:ff
fphana01:~ # _
    
```



iSCSI vNICs appear as ibft0 and ibft1 with their IP addresses assigned. Their interface configuration files ifcfg-eth2 and ifcfg-eth6 files are found in the /etc/sysconfig/network directory. We will be needing to create and configure the rest of the interface config files.

- c. Co-relating the outputs in step a and b above, we are able to determine the right IP address/network that need to should be assigned to the Ethernet interface. For the same an IP addressing scheme cheat sheet, as below, can be quite handy.

Host -Network	Inter-node	App Server	Admin	Backup
VLAN	220	223	76	224
Variable-info	<<var_internode_ipaddr>>	<<var_aapserver_ipaddr>>	<<var_mgmt_ipaddr>>	<<var_datasource_ipaddr>>
fphana01	192.168.220.210	192.168.223.210	192.168.76.210	192.168.224.210
Host -Network	HANA-data	Hana-log	iSCSI-A	iSCSI-B
VLAN	201	228	128	129
Variable-info	<<var_hana-data_ipaddr>>	<<var_hana-log_ipaddr>>	<<var_iscsi-a_ipaddr>>	<<var_iscsi-b_ipaddr>>lr>>
fphana01	192.168.201.210	192.168.228.210	192.168.128.50	192.168.129.51

- d. Assign the IP address and subnet mask for the ethernet interfaces based on the al the information we have so far.

```
#cd /etc/sysconfig/network
```

```
#vi ifcfg-eth0
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.220.210'
MTU='9000'
NAME='inter-node'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE=auto

#vi ifcfg-eth1
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.76.210'
MTU='1500'
NAME='admin'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE=auto

#vi ifcfg-eth3
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.223.210'
MTU='9000'
NAME='appserver'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE=auto

#vi ifcfg-eth4
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.201.210'
MTU='9000'
NAME='hana-data'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE=auto

#vi ifcfg-eth5
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.228.210'
MTU='9000'
NAME='hana-log'
NETMASK='255.255.255.0'
NETWORK=''
```

```

REMOTE_IPADDR=''
STARTMODE=auto

#vi ifcfg-eth7
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='192.168.224.210'
MTU='9000'
NAME='backup'
NETMASK='255.255.255.0'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE=auto

```

e. Add the default gateway.

```

#cd /etc/sysconfig/network
# vi routes
default <<var_mgmt_gateway_ip>> - -

```

f. Update the /etc/hosts with IP address of all networks and their alias hostnames:

```

fphana01:~ # vi /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

# special IPv6 addresses
::1           localhost ipv6-localhost ipv6-loopback
fe00::0       ipv6-localnet

ff00::0       ipv6-mcastprefix
ff02::1       ipv6-allnodes
ff02::2       ipv6-allrouters
ff02::3       ipv6-allhosts
#
## Internal Network
#
192.168.220.210 fphana01.ciscolab.local fphana01
#
## AppServer Network
#
192.168.223.210 fphana01a.ciscolab.local fphana01a
#
## Admin Network
#
192.168.76.210  fphana01m.ciscolab.local fphana01m

```

```
#
## Backup Network
#
192.168.224.210 fphana01b.ciscolab.local fphana01b
#
## HANA-data Network
#
192.168.201.210 fphana01d.ciscolab.local fphana01d
#
## HANA-log Network
#
192.168.228.210 fphana01l.ciscolab.local fphana01l
##
```

g. Create SWAP partition.

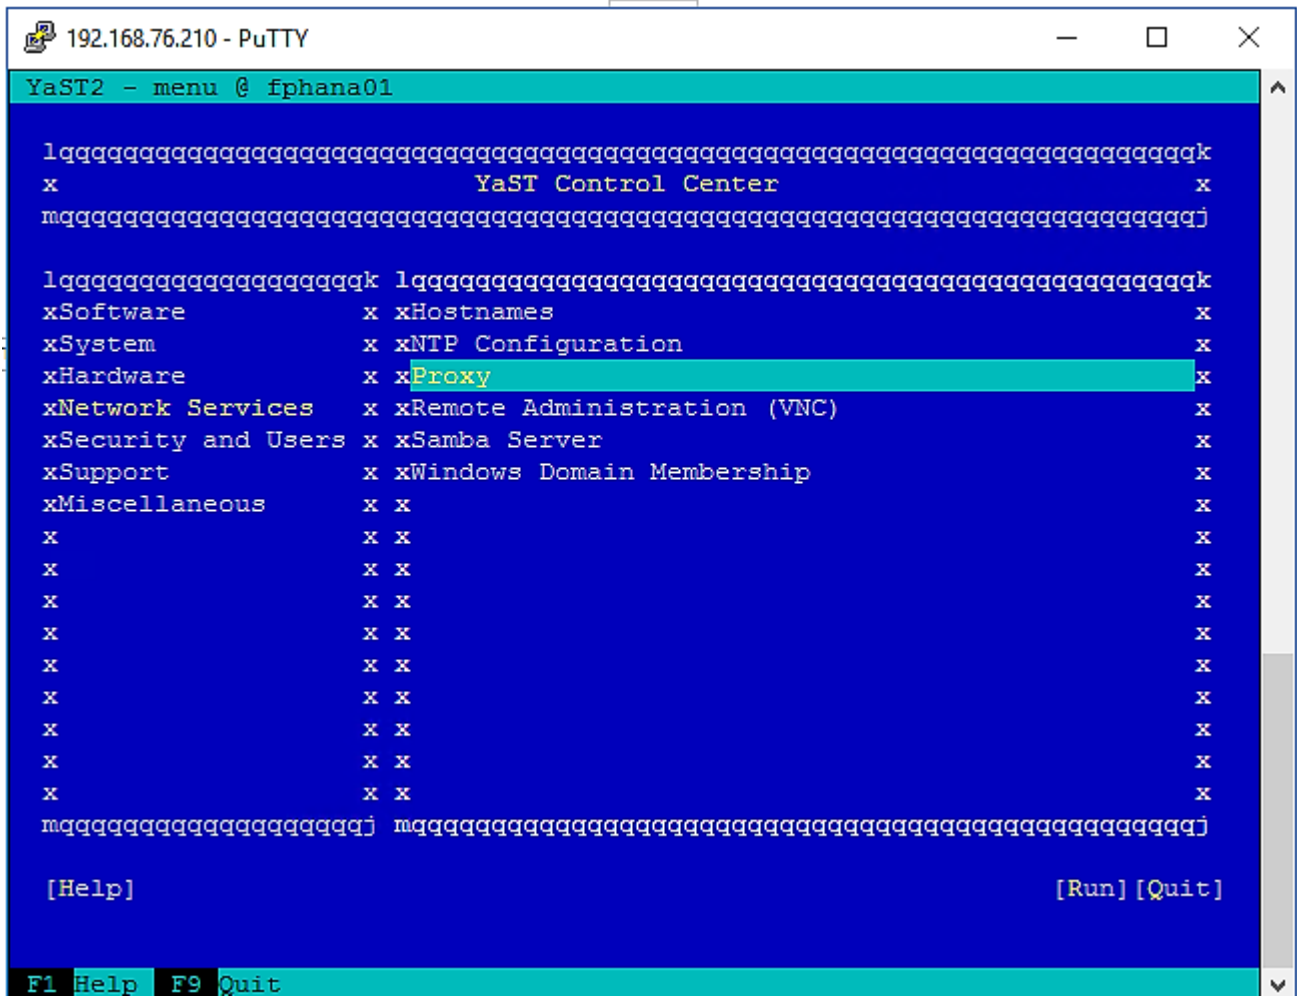
```
#dd if=/dev/zero of=/swap_01 bs=1024 count=2097152
#mkswap /swap_01
#chown root:disk /swap_01
#chmod 600 /swap_01
#swapon /swap_01
```

h. Update the /etc/fstab with swap filesystem information by appending this line.


```
/swap_01 swap swap defaults 0 0
```

8. Set up a proxy service, so that the appliance can reach the Internet.

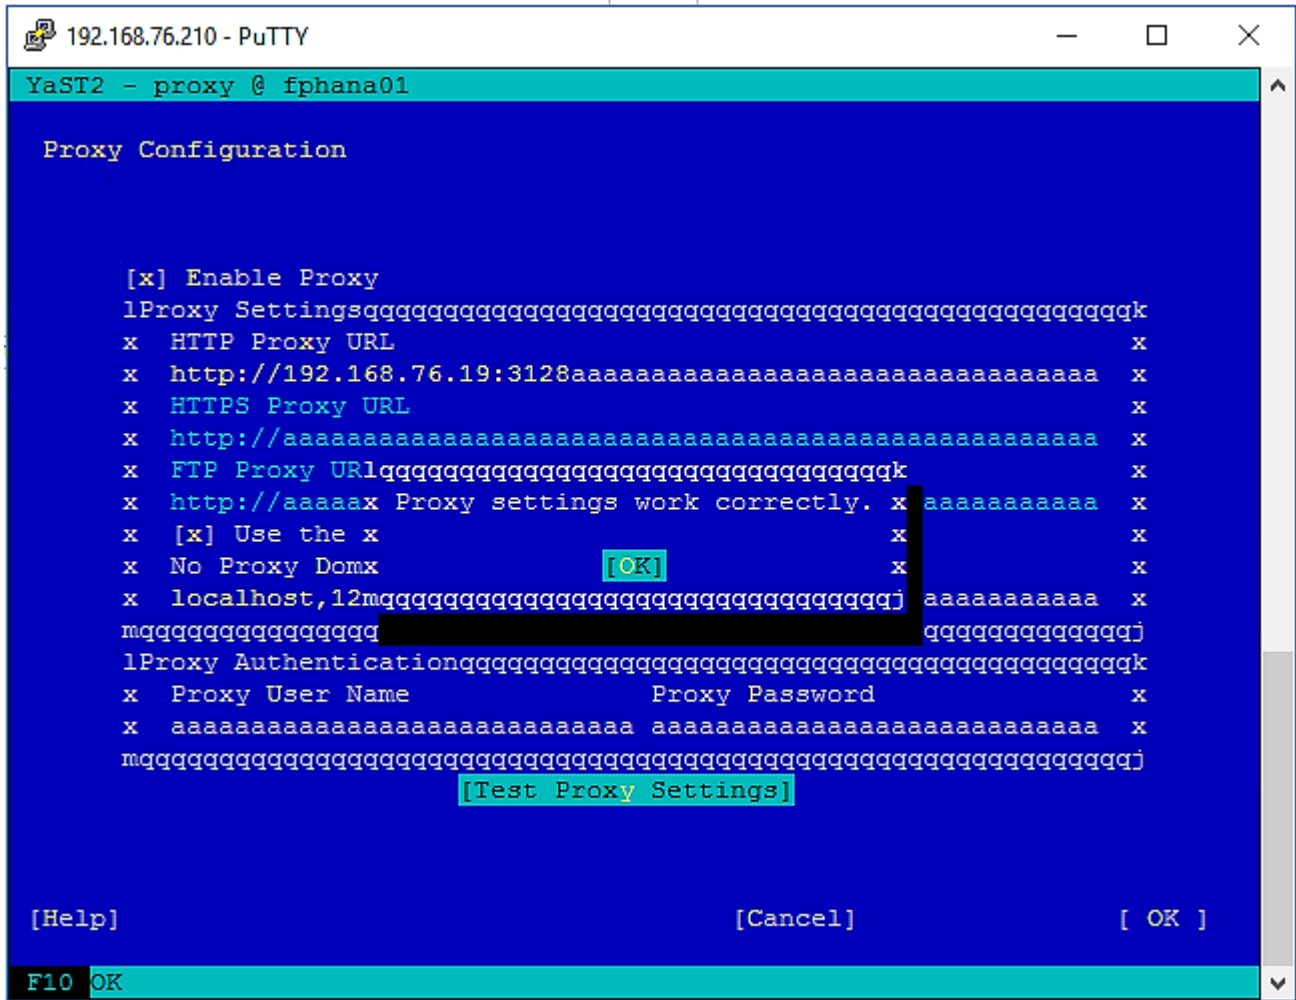
a. YaST2 – Key in the proxy server and port details. Select OK and then quit YaST to save the configuration.



- b. Select "Enable Proxy" > key in the <<proxy server IP address:port >> information and select "use same proxy for all Protocols" option.



c. Test the Proxy Settings to make they are working.



Reboot the system to effect the proxy server settings before going ahead with the registration step.

Update the SLES System

To update the SLES system, follow these steps:

1. Register the system with SUSE to get the latest patches. The system must have access to the Internet to proceed with this step.

```
#SUSEConnect -r <<registration_code>>
```

2. Update the system with the following command. Again, the system must have access to the Internet to proceed with this step.

```
#zypper update
```

3. Follow the on-screen instructions to complete the update process. Reboot the server and log in to the system again.

Set Kernel Parameters

Starting with SUSE Linux Enterprise Server12 SP1 the kernel parameter must be set in a configuration file in the /etc/sysctl.d directory. For example, a configuration file with the name 91-NetApp-HANA.conf must be created.

```
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.ipv4.tcp_rmem = 65536 16777216 16777216
net.ipv4.tcp_wmem = 65536 16777216 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

Install Cisco VIC Drivers

To install the Cisco VIC drivers, follow these steps:

1. Update fnic and enic drivers:
 - a. Based on the server type/model, processor version, OS release and version information download the Firmware bundle corresponding to the UCS Server firmware installed from the [Cisco UCS Hardware and Software Compatibility site](#)
 - b. Extract the rpm files of the fnic and enic drivers from the bundle over to the node.

```
fphana02:~ # cd /opt
fphana02:/opt # ll
total 3104
-rw-r--r-- 1 root root 1925500 Apr 23 15:41 cisco-enic-usnic-kmp-default-
3.2.272.23_k4.12.14_23-738.12.x86_64.rpm
-rw-r--r-- 1 root root 1240416 Apr 23 15:41 cisco-fnic-kmp-default-2.0.0.42-
77.0.x86_64.rpm

fphana02:/opt # rpm -ivh cisco-fnic-kmp-default-2.0.0.42-77.0.x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
 1:cisco-fnic-kmp-default-2.0.0.42_k##### [100%]

fphana02:/opt # rpm -ivh cisco-enic-usnic-kmp-default-3.2.272.23_k4.12.14_23-
738.12.x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
 1:cisco-enic-usnic-kmp-default-3.2.##### [100%]
Creating initrd: /boot/initrd-4.12.14-23-default
dracut: Executing: /usr/bin/dracut --logfile /var/log/YaST2/mkinitrd.log --force
/boot/initrd-4.12.14-23-default 4.12.14-23-default
dracut: dracut module 'dmraid' will not be installed, because command 'dmraid'
could not be found!
dracut: dracut module 'dmraid' will not be installed, because command 'dmraid'
could not be found!
dracut: *** Including module: bash ***
dracut: *** Including module: systemd ***
```

```

dracut: *** Including module: warpclock ***
dracut: *** Including module: systemd-initrd ***
dracut: *** Including module: il8n ***
dracut: *** Including module: network ***
dracut: *** Including module: drm ***
dracut: *** Including module: plymouth ***
dracut: *** Including module: kernel-modules ***
/usr/lib/dracut/modules.d/90kernel-modules/module-setup.sh: line 46: xhci-hcd:
command not found
dracut: *** Including module: kernel-network-modules ***
dracut: *** Including module: iscsi ***
dracut: *** Including module: rootfs-block ***
dracut: *** Including module: suse-btrfs ***
dracut: *** Including module: suse-xfs ***
dracut: *** Including module: terminfo ***
dracut: *** Including module: udev-rules ***
dracut: Skipping udev rule: 40-redhat.rules
dracut: Skipping udev rule: 50-firmware.rules
dracut: Skipping udev rule: 50-udev.rules
dracut: Skipping udev rule: 91-permissions.rules
dracut: Skipping udev rule: 80-drivers-modprobe.rules
dracut: *** Including module: dracut-systemd ***
dracut: *** Including module: haveged ***
dracut: *** Including module: usrmount ***
dracut: *** Including module: base ***
dracut: *** Including module: fs-lib ***
dracut: *** Including module: shutdown ***
dracut: *** Including module: suse ***
dracut: *** Including modules done ***
dracut: *** Installing kernel module dependencies and firmware ***
dracut: *** Installing kernel module dependencies and firmware done ***
dracut: *** Resolving executable dependencies ***
dracut: *** Resolving executable dependencies done***
dracut: *** Hardlinking files ***
dracut: *** Hardlinking files done ***
dracut: *** Stripping files ***
dracut: *** Stripping files done ***
dracut: *** Generating early-microcode cpio image ***
dracut: *** Constructing GenuineIntel.bin *****
dracut: *** Store current command line parameters ***
dracut: Stored kernel commandline:
dracut: rd.iscsi.ibft=1 rd.iscsi.firmware=1
dracut: root=UUID=1c6075f4-8c1d-481d-884d-9bd4a029e18e rootfstype=ext3
rootflags=rw,relatime,stripe=16,data=ordered
dracut: *** Creating image file '/boot/initrd-4.12.14-23-default' ***
dracut: *** Creating initramfs image file '/boot/initrd-4.12.14-23-default' done
***
fphana02:/opt #

```

2. Multipath configuration:

We disabled multipath during install time. It is recommended to install OS using a single path referencing the iSCSi device and enable the multipath configuration post OS installation and zypper update.

- a. At first check with `multipath -ll` may not return any result.
- b. Enable and start the multipath daemon and verify it active status.

```
fphana02:~ # multipath -ll
fphana02:~ # systemctl enable multipathd
Created symlink /etc/systemd/system/sysinit.target.wants/multipathd.service →
/usr/lib/systemd/system/multipathd.service.
Created symlink /etc/systemd/system/sockets.target.wants/multipathd.socket →
/usr/lib/systemd/system/multipathd.socket.

fphana02:~ # systemctl start multipathd

fphana02:~ # systemctl status multipathd
● multipathd.service - Device-Mapper Multipath Device Controller
   Loaded: loaded (/usr/lib/systemd/system/multipathd.service; enabled; vendor
  preset: disabled)
   Active: active (running) since Sun 2019-10-20 07:38:42 PDT; 34s ago
     Process: 32488 ExecStartPre=/sbin/modprobe -a scsi_dh_alua scsi_dh_emc
scsi_dh_rdac dm-multipath (code=exited, status=0/SUCCESS)
    Main PID: 32489 (multipathd)
       Status: "up"
         Tasks: 7
        CGroup: /system.slice/multipathd.service
                └─32489 /sbin/multipathd -d -s

Oct 20 07:38:42 fphana02 systemd[1]: Starting Device-Mapper Multipath Device
Controller...
Oct 20 07:38:42 fphana02 multipathd[32489]: -----start up-----
Oct 20 07:38:42 fphana02 multipathd[32489]: read /etc/multipath.conf
Oct 20 07:38:42 fphana02 multipathd[32489]: path checkers start up
Oct 20 07:38:42 fphana02 multipathd[32489]: 3618e728372e45ef022e8ccb60e0595ba: load
table [0 583983104 multipath 0 0 1 1 service-time 0 1 1 8:0 1]
Oct 20 07:38:42 fphana02 multipathd[32489]: 3600a0980383137384d244f3048386c49:
ignoring map
Oct 20 07:38:42 fphana02 multipathd[32489]: 3600a0980383137384d244f3048386c49:
ignoring map
Oct 20 07:38:42 fphana02 multipathd[32489]: 3618e728372e45ef022e8ccb60e0595ba:
event checker started
Oct 20 07:38:42 fphana02 systemd[1]: Started Device-Mapper Multipath Device
Controller.
```

- c. Check with `iscsiadm` commands to check for `iscsi` logins. If some paths are missing, discover the target and make sure the initiators are logged in to all ports.

```
fphana02:~ # iscsiadm -m session
tcp: [1] 192.168.129.12:3260,1030 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5 (non-flash)
tcp: [2] 192.168.128.12:3260,1027 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5 (non-flash)

fphana02:~ # iscsiadm -m discovery -t st -p 192.168.128.11
192.168.128.11:3260,1026 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5
192.168.129.12:3260,1030 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5
192.168.129.11:3260,1029 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5
192.168.128.12:3260,1027 iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5
```

```
fphana02:~ # iscsiadm -m node --loginall=all
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5, portal:
192.168.129.11,3260]
iscsiadm: default: 1 session requested, but 1 already present.
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5, portal:
192.168.128.11,3260]
iscsiadm: default: 1 session requested, but 1 already present.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5, portal:
192.168.129.11,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5, portal:
192.168.128.11,3260] successful.
iscsiadm: Could not log into all portals
```

- d. Restart multipathd and Enable and start the multipath daemon and verify it active status.

```
fphana02:~ # systemctl restart multipathd
fphana02:~ # multipath -r
fphana02:~ # multipath -ll
3600a0980383137384d244f3048386c49 dm-1 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 10:0:0:0 sde 8:64 active ready running
|  |- 9:0:0:0 sdd 8:48 active ready running
`--+- policy='service-time 0' prio=10 status=enabled
    |- 7:0:0:0 sdb 8:16 active ready running
    `-- 8:0:0:0 sdc 8:32 active ready running
```

- e. Create a Dracut configuration file exclusively for multipath.

```
#echo 'force_drivers+="dm_multipath dm_service_time"' >> /etc/dracut.conf.d/10-
mp.conf
```

- f. Rebuild the initrd and reboot the system.

```
fphana02:~ # dracut -f -a multipath
dracut: Executing: /usr/bin/dracut -f -a multipath
dracut: dracut module 'dmraid' will not be installed, because command 'dmraid'
could not be found!
dracut: dracut module 'dmraid' will not be installed, because command 'dmraid'
could not be found!
dracut: *** Including module: bash ***
dracut: *** Including module: systemd ***
dracut: *** Including module: warpclock ***
dracut: *** Including module: systemd-initrd ***
dracut: *** Including module: i18n ***
dracut: *** Including module: network ***
dracut: *** Including module: drm ***
dracut: *** Including module: plymouth ***
dracut: *** Including module: dm ***
dracut: Skipping udev rule: 64-device-mapper.rules
dracut: Skipping udev rule: 60-persistent-storage-dm.rules
dracut: Skipping udev rule: 55-dm.rules
```

```

dracut: *** Including module: kernel-modules ***
/usr/lib/dracut/modules.d/90kernel-modules/module-setup.sh: line 46: xhci-hcd:
command not found
dracut: *** Including module: kernel-network-modules ***
dracut: *** Including module: multipath ***
dracut: Skipping udev rule: 40-multipath.rules
dracut: *** Including module: iscsi ***
dracut: *** Including module: rootfs-block ***
dracut: *** Including module: suse-btrfs ***
dracut: *** Including module: suse-xfstools ***
dracut: *** Including module: terminfo ***
dracut: *** Including module: udev-rules ***
dracut: Skipping udev rule: 40-redhat.rules
dracut: Skipping udev rule: 50-firmware.rules
dracut: Skipping udev rule: 50-udev.rules
dracut: Skipping udev rule: 91-permissions.rules
dracut: Skipping udev rule: 80-drivers-modprobe.rules
dracut: *** Including module: dracut-systemd ***
dracut: *** Including module: haveged ***
dracut: *** Including module: usrmount ***
dracut: *** Including module: base ***
dracut: *** Including module: fs-lib ***
dracut: *** Including module: shutdown ***
dracut: *** Including module: suse ***
dracut: *** Including modules done ***
dracut: *** Installing kernel module dependencies and firmware ***
dracut: *** Installing kernel module dependencies and firmware done ***
dracut: *** Resolving executable dependencies ***
dracut: *** Resolving executable dependencies done***
dracut: *** Hardlinking files ***
dracut: *** Hardlinking files done ***
dracut: *** Stripping files ***
dracut: *** Stripping files done ***
dracut: *** Generating early-microcode cpio image ***
dracut: *** Constructing GenuineIntel.bin ****
dracut: *** Store current command line parameters ***
dracut: Stored kernel commandline:
dracut: rd.driver.pre=dm_multipath
rd.driver.pre=dm_service_time
dracut: rd.driver.pre=scsi_dh_alua rd.driver.pre=scsi_dh_emc
rd.driver.pre=scsi_dh_rdac rd.driver.pre=dm_multipath
dracut: rd.iscsi.ibft=1 rd.iscsi.firmware=1
dracut: root=/dev/disk/by-path/ip-192.168.128.12:3260-iscsi-iqn.1992-
08.com.netapp:sn.4f15529abe8611e9ae28d039ea00885a:vs.5-lun-0-part1 rootfstype=ext3
rootflags=rw,relatime,stripe=16,data=ordered
dracut: *** Creating image file '/boot/initrd-4.12.14-150.35-default' ***
dracut: *** Creating initramfs image file '/boot/initrd-4.12.14-150.35-default'
done ***

```

Implement SAP Notes recommendations

To optimize the HANA DB with SLES for SAP 15, follow the instructions in the [SAP Note 2684254](#):

1. SAP Note [1275776](#) describes how to apply recommended operating system settings for running SAP applications on SLES. There are three ways to implement the same – sapconf, saptune or manually. It is important to note When using sapconf or saptune, verify that parameters handled by these tools are not configured

elsewhere (e.g. boot parameter, sysctl.conf, and so on). This can cause inconsistent system behavior and makes debugging very hard.

2. This CVD uses the saptune (version 2) method which can prepare the operating system for SAP applications based on implementing specific SAP Notes
3. Install saptune:

```
#zypper install saptune
```

```
fphana01:~ # zypper in saptune
Refreshing service 'Basesystem_Module_15_x86_64'.
Refreshing service 'Desktop_Applications_Module_15_x86_64'.
Refreshing service 'Legacy_Module_15_x86_64'.
Refreshing service 'SAP_Applications_Module_15_x86_64'.
Refreshing service 'SUSE_Linux_Enterprise_High_Availability_Extension_15_x86_64'.
Refreshing service 'SUSE_Linux_Enterprise_Server_for_SAP_Applications_15_x86_64'.
Refreshing service 'Server_Applications_Module_15_x86_64'.
Loading repository data...
Reading installed packages...
'saptune' is already installed.
No update candidate for 'saptune-2.0.1-4.6.1.x86_64'. The highest available version is already installed.
Resolving package dependencies...

Nothing to do.
```

4. Configuration - activate saptune

```
#saptune daemon start
```

```
fphana01:~ # saptune daemon start
Starting daemon (tuned.service), this may take several seconds...
Daemon (tuned.service) has been enabled and started.
Your system has not yet been tuned. Please visit `saptune note` and `saptune solution` to start tuning.
fphana01:~ # █
```

5. All available solutions and notes can be listed with:

```
#saptune solution list
```

```
fphana01:~ # saptune solution list
All solutions (* denotes enabled solution, O denotes override file exists for solution, D denotes deprecated solutions):
BOBJ          - 941735 1771258 2578899 SAP_BOBJ
HANA          - 941735 1771258 1980196 2578899 2684254 2382421 2534844
D MAXDB       - 941735 1771258 2578899
NETWEAVER    - 941735 1771258 2578899
NETWEAVER+HANA - 941735 1771258 1980196 2578899 2684254 2382421 2534844
S4HANA-APP+DB - 941735 1771258 1980196 2578899 2684254 2382421 2534844
S4HANA-APPSERVER - 941735 1771258 2578899
S4HANA-DBSERVER - 941735 1771258 1980196 2578899 2684254 2382421 2534844
SAP-ASE      - 941735 1410736 1680803 1771258 2578899
fphana01:~ # █
```

6. Apply SAP HANA solution/notes:

```
#saptune solution apply HANA
```

```
fphana01:~ # saptune solution apply HANA
WARNING: [block] section detected: Traversing all block devices can take a considerable amount of time.
WARNING: Be aware: system-wide UserTasksMax is now set to infinity according to SAP recommendations.
This opens up entire system to fork-bomb style attacks.
All tuning options for the SAP solution have been applied successfully.
fphana01:~ # █
```

7. Verify the solution applied:

```
#saptune solution verify HANA
```

```
fphana01:~ # saptune solution verify HANA
WARNING: [block] section detected: Traversing all block devices can take a considerable amount of time.
```

SAPNote	Version	Parameter	Expected	Override	Actual	Compliant
1771258	5	LIMIT_@dba_hard_nofile	@dba hard nofile 65536		@dba hard nofile 65536	yes
1771258	5	LIMIT_@dba_soft_nofile	@dba soft nofile 65536		@dba soft nofile 65536	yes
1771258	5	LIMIT_@sapsys_hard_nofile	@sapsys hard nofile 65536		@sapsys hard nofile 65536	yes
1771258	5	LIMIT_@sapsys_soft_nofile	@sapsys soft nofile 65536		@sapsys soft nofile 65536	yes
1771258	5	LIMIT_@sdba_hard_nofile	@sdba hard nofile 65536		@sdba hard nofile 65536	yes
1771258	5	LIMIT_@sdba_soft_nofile	@sdba soft nofile 65536		@sdba soft nofile 65536	yes
1980196	7	vm.max_map_count	2147483647		2147483647	yes
2382421	33	net.core.somaxconn	4096		4096	yes
2382421	33	net.ipv4.tcp_max_syn_backlog	8192		8192	yes
2382421	33	net.ipv4.tcp_slow_start_after_idle	0		0	yes
2382421	33	net.ipv4.tcp_syn_retries	8		8	yes
2382421	33	net.ipv4.tcp_timestamps	1		1	yes
2382421	33	net.ipv4.tcp_window_scaling	1		1	yes
2534844	12	kernel.shmmax	32768		32768	yes
2578899	11	IO_SCHEDULER_sda	noop		noop	yes
2578899	11	IO_SCHEDULER_sdb	noop		noop	yes
2578899	11	IO_SCHEDULER_sdf	noop		noop	yes
2578899	11	UserTasksMax	infinity		infinity	yes
2578899	11	rpm:libopenssl1_0_0	1.0.2n-3.3.1		1.0.2p-3.22.1	yes [3]
2578899	11	rpm:libssh2-1	1.8.0-2.35		1.8.0-4.6.1	yes [3]
2578899	11	sysstat.service	start		start	yes
2578899	11	uuidd.socket	start		start	yes
2578899	11	vm.dirty_background_bytes	314572800		314572800	yes
2578899	11	vm.dirty_bytes	629145600		629145600	yes
2684254	5	KSM	0		0	yes
2684254	5	THP	never		never	yes
2684254	5	energy_perf_bias	all:0		all:0	yes
2684254	5	force_latency	70		10	yes
2684254	5	governor	all:performance		all:performance	yes
2684254	5	grub:intel_idle.max_cstate	1		NA	no [2] [3]
2684254	5	grub:numa_balancing	disable		NA	no [2] [3]
2684254	5	grub:processor.max_cstate	1		NA	no [2] [3]
2684254	5	grub:transparent_hugepage	never		NA	no [2] [3]
2684254	5	kernel.numa_balancing	0		0	yes
2684254	5	rpm:libopenssl1_0_0	1.0.2n-3.3.1		1.0.2p-3.22.1	yes [3]
2684254	5	rpm:libssh2-1	1.8.0-2.35		1.8.0-4.6.1	yes [3]
941735	11	ShmFileSystemSizeMB	1113685		1113685	yes
941735	11	VSZ_TMPFS_PERCENT	75		75	yes
941735	11	kernel.shmall	1152921504606846720		1152921504606846720	yes
941735	11	kernel.shmmax	18446744073709551615		18446744073709551615	yes

```
[2] setting is not available on the system
[3] value is only checked, but NOT set
```


System Provisioning for SAP HANA

This chapter describes the sequence of steps required to provision nodes for SAP HANA installation starting with the storage volume configuration, OS configuration needed to mount the storage volumes and subsequent use-case specific preparation tasks. The undelaying 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 12 shows the required variables used in this section.

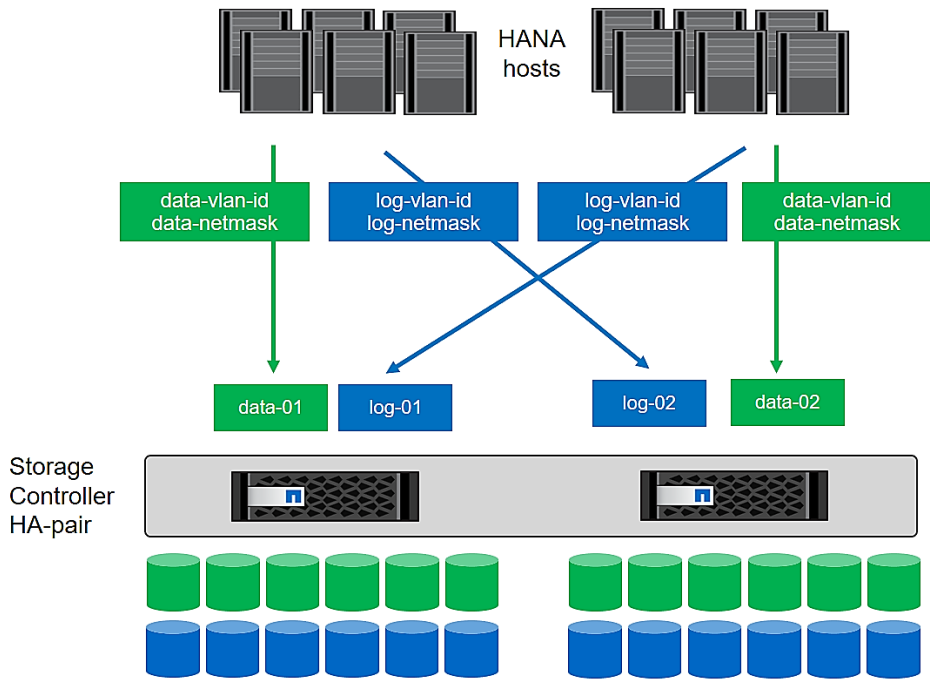
Table 12 Required Variables

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>	192.168.201.12
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 Figure 7, 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 7 Data and Log Volumes Distributed to the Storage Nodes



Configuring SAP HANA Single-Host Systems

Figure 8 illustrates 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 8 Volume Configuration for Four Single-Host SAP HANA Systems



Configure a data volume, a log volume, and a volume for /hana/shared for each SAP HANA host. Table 13 lists an example configuration for single-host SAP HANA systems.

Table 13 Volume Configuration for SAP HANA Single-Host Systems

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 at Controller b
Data, log, and shared volumes for system SID1	Data volume: SID1_data_mnt00001	Shared volume: SID1_shared		Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2		Log volume: SID2_log_mnt00001	Data volume: SID2_data_mnt00001	Shared volume: SID2_shared
Data, log, and shared volumes for system SID3	Shared volume: SID3_shared	Data volume: SID3_data_mnt00001	Log volume: SID3_log_mnt00001	
Data, log, and shared volumes for system SID4	Log volume: SID4_log_mnt00001		Shared volume: SID4_shared	Data volume: SID4_data_mnt00001

Table 14 lists 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 14 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	usr-sap	/usr/sap/SID
	shared	/hana/shared/SID

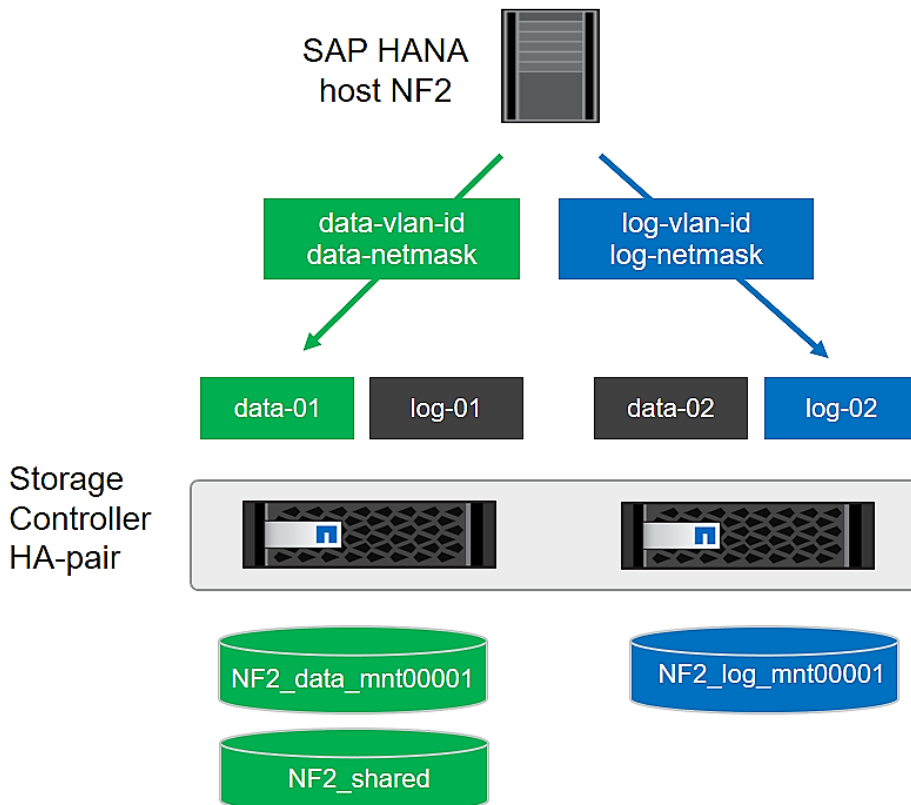
Configuration Example for a SAP HANA Single-Host System

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 9 shows the volumes that must be created on the storage nodes and the network paths used.

Figure 9 Required Volumes and Network Paths



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 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

```

Create Directories for HANA Shared Volume

To create the required directories for the hana shared volume mount the shared volume temporarily and create the required directories

```

lnx-jumphost:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
lnx-jumphost:/mnt # cd /mnt/tmp
lnx-jumphost:/mnt/tmp # mkdir shared usr-sap
lnx-jumphost:/mnt # cd ..
lnx-jumphost:/mnt/tmp # umount /mnt/tmp

```

Update the Load-Sharing Mirror Relation

To update the load-sharing mirror relation, run the following command:

```

snapmirror update-ls-set -source-path hana-svm:hana_rootvol

```

Create Mount Points on the node

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
```

Synchronize Domain information

To be able to mount the volumes inside the HANA nodes, the v4-id-domain information of NFS enabled hana-svm providing the HANA persistence access should tally with domain information in /etc/idmapd.conf of the HANA nodes.

On the NetApp command line fetch the information about the configured v4 domain as shown below:

```
sap-hana::> nfs show -vserver hana-svm -fields v4-id-domain
vserver  v4-id-domain
-----
hana-svm nfsv4domain.flexpod.com
sap-hana::>
```

Make sure the same is updated in the 'Domain' field on /etc/idmapd.conf file of HANA node.

```
[root@fprhel01 ~]# cat /etc/idmapd.conf
[General]
#Verbosity = 0
# The following should be set to the local NFSv4 domain name
# The default is the host's DNS domain name.
Domain = nfsv4domain.flexpod.com
```

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 15 lists the required mount options.

This example uses NFSv4.1 to connect the storage. However, NFSv3 is supported for SAP HANA single host systems.

For NFSv3, 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 connections to the storage system over 10GbE, you must set the transfer size to 1MB to achieve the expected throughput values.

Table 15 Mount Options

Common Parameter	NFSv4.1	NFS Transfer Size with ONTAP 9
rw, bg, hard, timeo=600, intr, noatime	vers=4, minorversion=1, lock	rsize=1048576, wsize=1048576

To mount the file systems during system boot using the `/etc/fstab` configuration file, follow these steps:



The following examples show an SAP HANA database with SID=NF2 using NFSv4.1 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=4,minorversion=1,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,lock 0 0
<node02-log_lif01-ip>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,bg,vers=4,minorversion=1,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,lock 0 0
<node01-data_lif01-ip>:/NF2_shared/usr-sap /usr/sap/NF2 nfs
rw,bg,vers=4,minorversion=1,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,lock 0 0
<node01-data_lif01-ip>:/NF2_shared/shared /hana/shared nfs
rw,bg,vers=4,minorversion=1,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,lock 0 0
```

2. Run `mount -a` to mount the file systems on the host.

Persistent Memory Configuration

For detailed information, see: [Cisco UCS for SAP HANA with Intel Optane DC Persistent Memory Module white paper](#)

This section addresses the use-case of Single-host system implementation on a node configured with Intel Persistent Memory and extends configuration scope to include the same.

The utility `ipmctl` is used for configuring and managing Intel Optane DC persistent memory modules (DCPMM) and the `ndctl` utility library is required for managing the `libnvdimm` (non-volatile memory device) sub-system in the Linux kernel.

To configure the persistent memory, follow these steps:

1. `ssh` to the Server as root.
2. Install the `ipmctl` host utility

```
# zypper in ipmctl
The following 2 NEW packages are going to be installed:
  ipmctl libndctl6
2 new packages to install.
Overall download size: 487.7 KiB. Already cached: 0 B. After the operation, additional 3.4
MiB will be used.
Continue? [y/n/v/...? shows all options] (y): y
Retrieving package libndctl6-63-3.5.1.x86_64
(1/2), 87.6 KiB (188.0 KiB unpacked)
Retrieving: libndctl6-63-3.5.1.x86_64.rpm .....[done]
Retrieving package ipmctl-01.00.00.3440-1.6.1.x86_64
(2/2), 400.1 KiB ( 3.2 MiB unpacked)
Retrieving: ipmctl-01.00.00.3440-1.6.1.x86_64.rpm .....[done]
```

```

Checking for file conflicts: .....[done]
(1/2) Installing: libndctl6-63-3.5.1.x86_64.....[done]
(2/2) Installing: ipmctl-01.00.00.3440-1.6.1.x86_64.....[done]

```

3. Install the ndctl utility library

```

# zypper in ndctl
The following NEW package is going to be installed:
ndctl
1 new package to install.
Overall download size: 147.2 KiB. Already cached: 0 B. After the operation, additional 252.1
KiB will be used.
Continue? [y/n/v/...? shows all options] (y): y
Retrieving package ndctl-63-3.5.1.x86_64
(1/1), 147.2 KiB (252.1 KiB unpacked)
Retrieving: ndctl-63-3.5.1.x86_64.rpm .....[done]
Checking for file conflicts: .....[done]
(1/1) Installing: ndctl-63-3.5.1.x86_64 .....[done]

```

4. Confirm the persistent memory modules are discovered in the system and verify the software can communicate with them.

```

# ipmctl show -dimms

```

DimmID	Capacity	HealthState	ActionRequired	LockState	FWVersion
0x0001	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x0011	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x0021	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x0101	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x0111	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x0121	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1001	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1011	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1021	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1101	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1111	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x1121	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2001	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2011	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2021	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2101	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2111	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x2121	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3001	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3011	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3021	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3101	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3111	252.4 GiB	Healthy	0	Disabled	01.02.00.5375
0x3121	252.4 GiB	Healthy	0	Disabled	01.02.00.5375

5. Create the goal

```

# ipmctl create -goal

```

The following configuration will be applied:

SocketID	DimmID	MemorySize	AppDirect1Size	AppDirect2Size
0x0000	0x0001	0.0 GiB	252.0 GiB	0.0 GiB
0x0000	0x0011	0.0 GiB	252.0 GiB	0.0 GiB
0x0000	0x0021	0.0 GiB	252.0 GiB	0.0 GiB


```

0x0000 | 0x0101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
Do you want to continue? [y/n] y
Created following region configuration goal
SocketID | DimmID | MemorySize | AppDirect1Size | AppDirect2Size
=====
0x0000 | 0x0001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0000 | 0x0121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0001 | 0x1121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0002 | 0x2121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3001 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3011 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3021 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3101 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3111 | 0.0 GiB | 252.0 GiB | 0.0 GiB
0x0003 | 0x3121 | 0.0 GiB | 252.0 GiB | 0.0 GiB
A reboot is required to process new memory allocation goals.

```

6. Reboot the server for the new memory allocations.

7. Verify the regions created

```

# ipmctl show -region

SocketID | ISetID | PersistentMemoryType | Capacity | FreeCapacity |
HealthState
=====
===

```

0x0000	0x96467f486ebf2ccc	AppDirect	1512.0 GiB	1512.0 GiB	Healthy
0x0001	0x163e7f48a6eb2ccc	AppDirect	1512.0 GiB	1512.0 GiB	Healthy
0x0002	0xd8787f48c4af2ccc	AppDirect	1512.0 GiB	1512.0 GiB	Healthy
0x0003	0x1ac47f482bb32ccc	AppDirect	1512.0 GiB	1512.0 GiB	Healthy

8. Create a name space for each region; on a server with a total of 4 CPU invoke the command four times.

```
ndctl create-namespace
```

9. List the active name spaces created before

```
# ndctl list
[
  {
    "dev": "namespace24.0",
    "mode": "fsdax",
    "map": "dev",
    "size": 1598128390144,
    "uuid": "60f803c0-d6f6-4a0c-b522-393e74d25279",
    "blockdev": "pmem24"
  },
  {
    "dev": "namespace26.0",
    "mode": "fsdax",
    "map": "dev",
    "size": 1598128390144,
    "uuid": "306b6e4b-f2e0-49b0-95e8-0d44602c2204",
    "blockdev": "pmem26"
  },
  {
    "dev": "namespace25.0",
    "mode": "fsdax",
    "map": "dev",
    "size": 1598128390144,
    "uuid": "b08acc26-a196-4de1-ae1c-b088058410ee",
    "blockdev": "pmem25"
  },
  {
    "dev": "namespace27.0",
    "mode": "fsdax",
    "map": "dev",
    "size": 1598128390144,
    "uuid": "1a20712f-cce0-49e9-b871-3b424c740ff4",
    "blockdev": "pmem27"
  }
]
```

10. Create a file system and mount the persistent memory modules

```
mkfs -t xfs -f /dev/pmem24
mkfs -t xfs -f /dev/pmem25
mkfs -t xfs -f /dev/pmem26
mkfs -t xfs -f /dev/pmem27

mkdir -p /hana/pmem/nvmem0
mkdir -p /hana/pmem/nvmem1
mkdir -p /hana/pmem/nvmem2
mkdir -p /hana/pmem/nvmem3

mount -t xfs -o dax /dev/pmem24 /hana/pmem/nvmem0
mount -t xfs -o dax /dev/pmem25 /hana/pmem/nvmem1
```

```
mount -t xfs -o dax /dev/pmem26 /hana/pmem/nvmem2
mount -t xfs -o dax /dev/pmem26 /hana/pmem/nvmem3
```

11. Add the mount points to the `/etc/fstab` file to make them permanent

```
# vi /etc/fstab
/dev/pmem24 on /hana/pmem/nvmem0 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem25 on /hana/pmem/nvmem1 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem26 on /hana/pmem/nvmem2 type xfs (rw,relatime,attr2,dax,inode64,noquota)
/dev/pmem37 on /hana/pmem/nvmem3 type xfs (rw,relatime,attr2,dax,inode64,noquota)
```

Configure the Base Path to Use Persistent Memory

Post the SAP HANA Single-host system install the directory that SAP HANA uses as its base path must point to the XFS file system. Define the base path location with the configuration parameter `basepath_persistent_memory_volumes` in the `[persistence]` section of the SAP HANA `global.ini` file. This section can contain multiple locations separated by semicolons.


Changes to this parameter require a restart of SAP HANA services.

```
[persistence]

basepath_datavolumes = /hana/data/<SID>
basepath_logvolumes = /hana/log/<SID>
basepath_persistent_memory_volumes =
/hana/pmem/nvmem0;/hana/pmem/nvmem1;/hana/pmem/nvmem2;/hana/pmem/nvmem3
```

Configure SAP HANA Multiple-Host Systems

The figure below 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.



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 16 shows an example configuration for a multiple-host SAP HANA system with four active hosts.

Table 16 Volume Configuration for SAP HANA Multiple-Host Systems

Purpose	Aggregate 1 at Controller A	Aggregate 2 at Controller A	Aggregate 1 at Controller B	Aggregate 2 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	N/A	N/A

Table 17 lists 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 17 Mount Points for Multiple-Host Systems

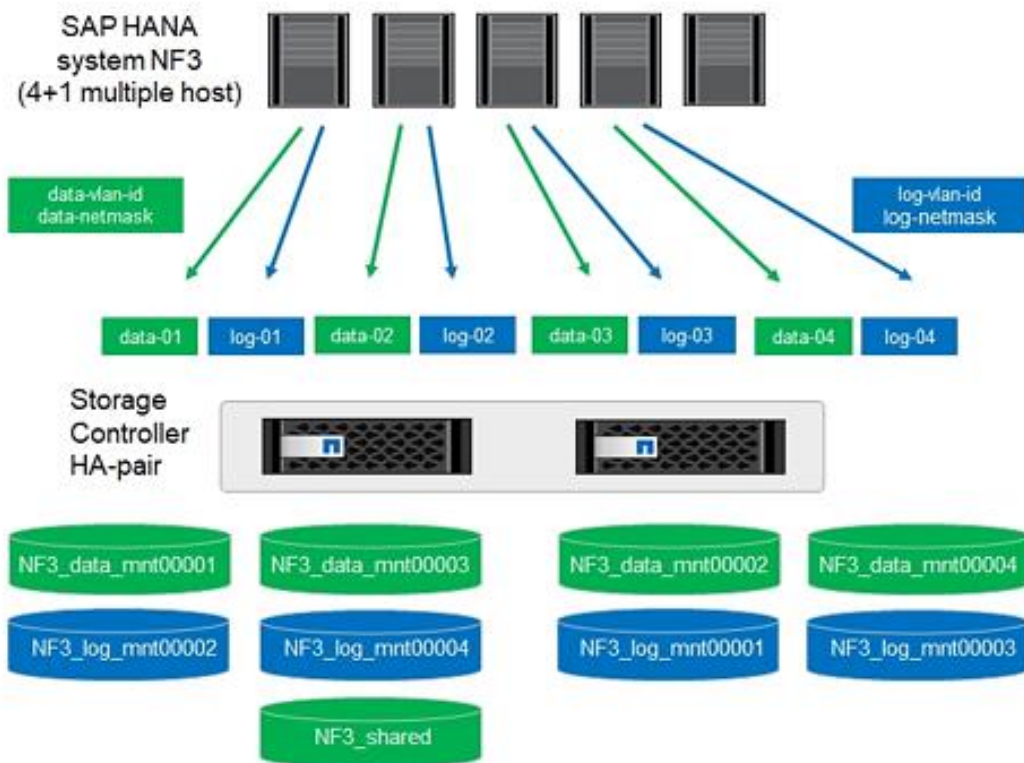
Junction Path	Directory	Mount Point at SAP HANA Host	Note
SID_data_mnt00001		/hana/data/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001		/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002		/hana/data/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002		/hana/log/SID/mnt00002	Mounted at all 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 4+1 SAP HANA multiple-host database with SID=NF3 and a server with a RAM size of 2TB. 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.

The figure below shows the volumes that must be created on the storage nodes and the network paths used.



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 aggr1_1 -size 2500GB -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 aggr1_2 -size 2500GB -state online
-junction-path /NF3_data_mnt00002 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -
space-guarantee none
volume create -vserver hana-svm -volume NF3_data_mnt00003 -aggregate aggr2_1 -size 2500GB -state online
-junction-path /NF3_data_mnt00003 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -
space-guarantee none
volume create -vserver hana-svm -volume NF3_data_mnt00004 -aggregate aggr2_2 -size 2500GB -state online
-junction-path /NF3_data_mnt00004 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -
space-guarantee none

vol modify -vserver hana-svm -volume NF3_data_mnt0000* -snapdir-access false
set advanced
vol modify -vserver hana-svm -volume NF3_data_mnt0000* -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 aggr1_2 -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 aggr1_1 -size 512GB -state online -
junction-path /NF3_log_mnt00002 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none
volume create -vserver hana-svm -volume NF3_log_mnt00003 -aggregate aggr2_2 -size 512GB -state online -
junction-path /NF3_log_mnt00003 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none
```

```

volume create -vsserver hana-svm -volume NF3_log_mnt00004 -aggregate aggr2_1 -size 512GB -state online -
junction-path /NF3_log_mnt00004 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vsserver hana-svm -volume NF3_log_mnt0000* -snapdir-access false
set advanced
vol modify -vsserver hana-svm -volume NF3_log_mnt0000* -atime-update false
set admin

```

Create HANA Shared Volume and Adjust Volume Options

To create a HANA shared volume and adjust the volume options, run the following commands:

```

volume create -vsserver hana-svm -volume NF3_shared -aggregate aggr1_1 -size 2500GB -state online -
junction-path /NF3_shared -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vsserver hana-svm -volume NF3_shared -snapdir-access false
set advanced
vol modify -vsserver hana-svm -volume NF3_shared -atime-update false
set admin

```

Create Directories of HANA Shared Volume

To create the required directories for the HANA shared volume, mount the shared volume temporarily and create the required directories

```

lnx-jumphost:/mnt # mount <storage-hostname>:/NF3_shared /mnt/tmp
lnx-jumphost:/mnt # cd /mnt/tmp
lnx-jumphost:/mnt/tmp # mkdir shared
lnx-jumphost:/mnt/tmp # mkdir usr-sap-host1
lnx-jumphost:/mnt/tmp # mkdir usr-sap-host2
lnx-jumphost:/mnt/tmp # mkdir usr-sap-host3
lnx-jumphost:/mnt/tmp # mkdir usr-sap-host4
lnx-jumphost:/mnt/tmp # mkdir usr-sap-host5
lnx-jumphost:/mnt # cd ..
lnx-jumphost:/mnt/tmp # umount /mnt/tmp

```

Update Load-Sharing Mirror Relation

To update the load-sharing mirror relation, run the following command:

```

snapmirror update-ls-set -source-path hana-svm:hana_rootvol

```

Create Mount Points

For a multiple-host system, create mount points and set the permissions on all worker and standby hosts as follows.

1. Create mount points for the first worker host.

```

mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/data/NF3/mnt00003
mkdir -p /hana/data/NF3/mnt00004
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00003
mkdir -p /hana/log/NF3/mnt00004

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/data/NF3/mnt00003
mkdir -p /hana/data/NF3/mnt00004
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00003
mkdir -p /hana/log/NF3/mnt00004
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 third worker host.

```

mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/data/NF3/mnt00003
mkdir -p /hana/data/NF3/mnt00004
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00003
mkdir -p /hana/log/NF3/mnt00004
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

```

4. Create mount points for the fourth worker host.

```

mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/data/NF3/mnt00003
mkdir -p /hana/data/NF3/mnt00004
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00003
mkdir -p /hana/log/NF3/mnt00004
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

```

5. Create mount points for the standby host.

```

mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/data/NF3/mnt00003
mkdir -p /hana/data/NF3/mnt00004
mkdir -p /hana/log/NF3/mnt00001

```



```

mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00003
mkdir -p /hana/log/NF3/mnt00004
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

```

Synchronize Domain information

To be able to mount the volumes inside the HANA nodes, the v4-id-domain information of NFS enabled hana-svm providing the HANA persistence access should tally with domain information in /etc/idmapd.conf of the HANA nodes.

On the NetApp command line fetch the information about the configured v4 domain as below:

```

sap-hana::> nfs show -vserver hana-svm -fields v4-id-domain
vserver  v4-id-domain
-----
hana-svm nfsv4domain.flexpod.com

sap-hana::> █

```

Make sure the same is updated in the 'Domain' filed on /etc/idmapd.conf file of HANA node.

```

fphana01:/hana/shared/tools # cat /etc/idmapd.conf
[General]

Verbosity = 0
Pipefs-Directory = /var/lib/nfs/rpc_pipefs
Domain = nfsv4domain.flexpod.com

[Mapping]

Nobody-User = nobody
Nobody-Group = nobody

[Translation]
Method=nsswitch
fphana01:/hana/shared/tools # █

```

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/data/NF3/mnt00003
- /hana/data/NF3/mnt00004
- /hana/log/NF3/mnt00001

- /hana/log/NF3/mnt00002
- /hana/log/NF3/mnt00003
- /hana/log/NF3/mnt00004
- /hana/shared
- /usr/sap/NF3

Table 18 lists 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 18 Mount Options

Common Parameter	NFSv4.1	NFS Transfer Size with ONTAP 9
rw,bg,hard, timeo=600,intr, noatime	vers=4,minorversion=1,lock	rsize=1048576, wsize=1048576

The following examples show a SAP HANA database with SID=NF3 using NFSv4.1 and an NFS transfer size of 1MB. To mount the file systems during system boot using the /etc/fstab configuration file, follow these 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=4,minorversion=1,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=4,minorversion=1,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node01-data_lif01-ip>:/NF3_data_mnt00003 /hana/data/NF3/mnt00003 nfs
rw,bg,vers=4,minorversion=1,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node02-data_lif01-ip>:/NF3_data_mnt00004 /hana/data/NF3/mnt00004 nfs
rw,bg,vers=4,minorversion=1,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=4,minorversion=1,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=4,minorversion=1,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node02-log_lif01-ip>:/NF3_log_mnt00003 /hana/log/NF3/mnt00003 nfs
rw,bg,vers=4,minorversion=1,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node01-log_lif01-ip>:/NF3_log_mnt00004 /hana/log/NF3/mnt00004 nfs
rw,bg,vers=4,minorversion=1,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=4,minorversion=1,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=4,minorversion=1,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
```

2. Run `mount -a` on each host to mount the file systems.
3. For scale-out system, all nodes should be able to resolve Internal network IP address. Below is an example of 4 node scale-out system host file with all the network defined in the /etc/hosts file:

```

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

# special IPv6 addresses
::1           localhost ipv6-localhost ipv6-loopback

fe00::0       ipv6-localnet

ff00::0       ipv6-mcastprefix
ff02::1       ipv6-allnodes
ff02::2       ipv6-allrouters
ff02::3       ipv6-allhosts
192.168.201.11 hana-data-a
192.168.201.12 hana-data-b
192.168.228.11 hana-log-a
192.168.228.12 hana-log-b
#
## Inter-node Network
#
192.168.220.210 fphana01.ciscolab.local fphana01
192.168.220.211 fphana02.ciscolab.local fphana02
192.168.220.212 fphana03.ciscolab.local fphana03
192.168.220.213 fphana04.ciscolab.local fphana04
#
## Storage Data Network
#
192.168.201.210 fphana01d.ciscolab.local fphana01d
192.168.201.211 fphana02d.ciscolab.local fphana02d
192.168.201.212 fphana03d.ciscolab.local fphana03d
192.168.201.213 fphana04d.ciscolab.local fphana04d
#
## Storage Log Network
#
192.168.228.210 fphana01l.ciscolab.local fphana01l
192.168.228.211 fphana02l.ciscolab.local fphana02l
192.168.228.212 fphana03l.ciscolab.local fphana03l
192.168.228.213 fphana04l.ciscolab.local fphana04l
#
## Client Network
#
192.168.222.210 fphana01c.ciscolab.local fphana01c
192.168.222.211 fphana02c.ciscolab.local fphana02c
192.168.222.212 fphana03c.ciscolab.local fphana03c
192.168.222.213 fphana04c.ciscolab.local fphana04c
#
## AppServer Network
#
192.168.223.210 fphana01a.ciscolab.local fphana01a
192.168.223.211 fphana02a.ciscolab.local fphana02a
192.168.223.212 fphana03a.ciscolab.local fphana03a
192.168.223.213 fphana04a.ciscolab.local fphana04a
#
## Admin Network
#
192.168.76.210  fphana01m.ciscolab.local fphana01m
192.168.76.211  fphana02m.ciscolab.local fphana02m
192.168.76.212  fphana03m.ciscolab.local fphana03m
192.168.76.213  fphana04m.ciscolab.local fphana04m
#

```

```
#
## Backup Network
#
192.168.224.210 fphana01b.ciscolab.local fphana01b
192.168.224.211 fphana02b.ciscolab.local fphana02b
192.168.224.212 fphana03b.ciscolab.local fphana03b
192.168.224.213 fphana04b.ciscolab.local fphana04b
#
## Replication Network
#
192.168.225.210 fphana01r.ciscolab.local fphana01r
192.168.225.211 fphana02r.ciscolab.local fphana02r
192.168.225.212 fphana03r.ciscolab.local fphana03r
192.168.225.213 fphana04r.ciscolab.local fphana04r
```

Passwordless Authentication

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. 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 fphana02"
```

```
ssh-copy-id -i /root/.ssh/id_rsa.pub fphana02
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.
```

3. Repeat steps 1-2 for all the servers in the scale-out system so that each node has the keys copied over from all other nodes.

A sample global.ini file from the installed multi-host SAP HANA system is as below:

```
[communication]
listeninterface = .global

[multidb]
mode = multidb
database_isolation = low
singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/NF3
basepath_logvolumes = /hana/log/NF3
```

SAP HANA Installation

For information about the 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
- [2578899 - SUSE Linux Enterprise Server 15: Installation Note](#)
- [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 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 Technical reports

[TR-4435-SAP HANA on NetApp AFF Systems with NFS](#)

[TR-3580-NFSv4 Enhancements and Best Practices](#)

[TR-4614-SAP HANA Backup and Recovery with SnapCenter](#)

[TR-4646-SAP HANA Disaster Recovery with Asynchronous Storage Replication](#)

High-Availability Configuration for Multiple-host

Since you are using NFSv4.1, a specific HA configuration for multiple hosts SAP HANA databases is not necessary.

NFSv4 locking ensures that only one hosts can access the SAP HANA data and log files at a time.

SAP HANA Installation Preparations for NFSv4

NFS version 4 and higher requires user authentication. This authentication can be accomplished by using a central user management tool such as an LDAP server or local user accounts. The following sections describe how to configure local user accounts.

The administration user <sidadm> and the sapsys group must be created manually on the SAP HANA hosts and the storage controllers before the installation of SAP HANA software begins.

SAP HANA Hosts

If the sapsys group doesn't exist, you must create it on the SAP HANA host. You must choose a unique group ID that does not conflict with the existing group IDs on the storage controllers.

Create the user <sidadm> on the SAP HANA host. You must choose a unique ID that does not conflict with existing user IDs on the storage controllers.

For a multiple-host SAP HANA system, the user and group ID must be the same on all SAP HANA hosts. The group and user are created on the other SAP HANA hosts by copying the affected lines in `/etc/group` and `/etc/passwd` from the source system to all other SAP HANA hosts.



The NFSv4 domain must be set to the same value on all Linux servers (`/etc/idmapd.conf`) and SVMs. Set the domain parameter "Domain = < nfsv4domain.flexpod.com >" in the file `/etc/idmapd.conf` for the Linux hosts.

Storage Controllers

The user ID and group ID must be the same on the SAP HANA hosts and the storage controllers. To create the group and user, run the following commands on the storage cluster:

```
vserver services unix-group create -vserver <vserver> -name <group name> -id <group id>
vserver services unix-user create -vserver <vserver> -user <user name> -id <user-id> -primary-gid <group id>
```


SAP HANA Data Volume Size

A default SAP HANA instance uses only one data volume per SAP HANA service. Due to the max file size limitation of the file system, NetApp recommends limiting max data volume size.

To do so automatically set the following parameter in the global.ini within section [persistence]:

```
datavolume_striping = true  
datavolume_striping_size_gb = 8000
```

This created a new data volume when the limit of 8000GB is reached.

SAP Note 240005 question 15 provides more information: <https://launchpad.support.sap.com/#/notes/2400005>

Monitor SAP HANA with AppDynamics

Introduction

AppDynamics is an Application Performance Monitoring (APM) Platform that helps you to understand and optimize the performance of your business, from its software to infrastructure to business journeys.

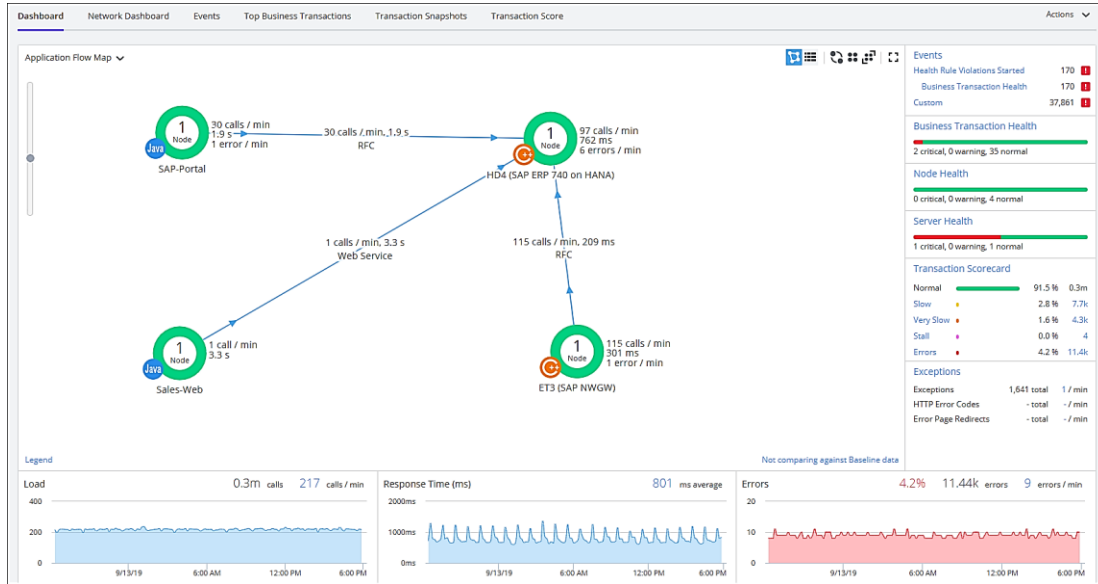
The AppDynamics APM Platform enables you to monitor and manage your entire application-delivery ecosystem, from the mobile app or browser client request through your network, backend databases and application servers and more. AppDynamics APM gives you a single view across your application landscape, letting you quickly navigate from the global perspective of your distributed application right down to the call graphs or exception reports generated on individual hosts.

AppDynamics has an agent-based architecture. Once our agents are installed it gives you a dynamic flow map or topography of your application. It uses the concept of traffic lights to indicate the health of your application (green is good, yellow is slow and red indicates potential issues) with dynamic baselining. AppDynamics measures application performance based on business transactions which essentially are the key functionality of the application. When the application deviates from the baseline AppDynamics captures and provides deeper diagnostic information to help be more proactive in troubleshooting and reduce the MTTR (mean time to resolution).

SAP Landscape Monitoring

AppDynamics has a one of its kind ABAP agent for monitoring SAP ABAP systems. We have comprehensive coverage of the SAP landscape with our ABAP, Java, .net and Server visibility agents. In addition, Datavard Insights extends the AppDynamics for SAP solution with system-level monitoring for the overall SAP systems and SAP HANA databases. While AppDynamics agents provides transaction-level visibility, Datavard Insights collects performance metrics, logs and events, including processes outside of the user business transactions, such as background jobs or IDocs processing.

The complex and proprietary nature of SAP applications makes it difficult to diagnose issues. AppDynamics allows enterprises to instrument SAP applications, monitor performance, and understand the root cause of performance bottlenecks.



Trial Registration

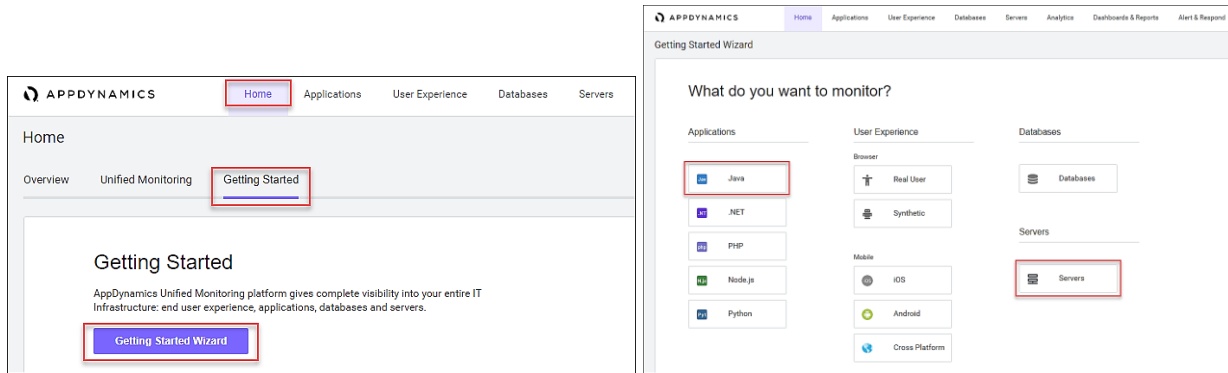
To register, follow these steps:

1. Connect to <https://www.appdynamics.com/free-trial/>
2. Provide the details to sign up for a free trial utilizing an AppDynamics SaaS controller.

The screenshot shows the AppDynamics sign-up page with the following content:


- Header:** "Get started quickly with AppDynamics" and "Start your free 15-day SaaS trial now, easily, with no credit card required".
- Benefits:**
 - Application Performance Management:** Monitor and manage end-to-end performance of complex distributed applications.
 - End User Monitoring:** Deliver a better user experience with Mobile and Browser Real-User Monitoring and Browser Synthetic Monitoring.
 - Infrastructure Visibility:** Unlock deeper insights by correlating server and database performance with application performance.
- Sign up for a free SaaS trial form:**
 - Fields: First Name, Last Name, Company Email, AppDynamics, Job Title, Phone Number, Create Password.
 - Checkbox: "I agree to the Evaluation Terms and Privacy Policy."
 - Button: "CREATE A FREE ACCOUNT"
 - Footer: "Already have an account? Sign In." and "*For On-Premise trials, please contact us."

3. Once the AppDynamics SaaS Controller has been provisioned, you will receive an email with the information you need for you to login to the Controller.
4. You can download the Java Agent and the Server Visibility Agent directly from the Controller



5. You can use the email and password you provided to sign up for the trial to login to the agent download site at the URL listed below and download the ABAP Agent:

<https://download.appdynamics.com>



Agent Installation

AppDynamics has several types of agents to monitor different language applications to user Experience to Infrastructure monitoring. Based on the SAP landscape and the underlying technology of the SAP systems the agents are installed.

The most frequently installed agents are:

1. Java Agent - For Java based SAP Systems
2. ABAP Agent - For ABAP based SAP systems
3. Server Visibility Agent - Provides extended hardware metrics and Service Availability Monitoring

Prerequisites

Please see the link below to verify the supported SAP environments:

<https://docs.appdynamics.com/display/SAP/SAP+Supported+Environments>

Java Agent Installation

The Java Agent must be installed on SAP Java application servers (e.g. Enterprise Portal and PO application servers).

The high-level steps for installing the Java Agent are listed below:

1. Ensure you are logged into the host server as the appropriate <SID>adm OS user
2. Create the permanent installation directory for the agent
3. Unzip the file in the permanent directory (i.e. */usr/sap/appdyn/app*)
4. Change the permissions on the contents of the agents' installation directory to give full read/write privileges at the owner and group levels
5. Edit the configuration file in the agent installation directory to configure the agent to connect to the controller, provide the identity of the JVM, and so on.
6. You will need to add parameters to the SAP JVM to start the Java agent when the SAP system is started up by logging into the SAP app server as the "sidadm" user
7. Use the SAP NetWeaver Administrator or the AS Java Config Tool (depending on your SAP system) to edit the JVM startup parameters. For more detailed information, see the link below:

[Configuring AppDynamics Java Agent in SAP](#)

8. Restart the SAP JVM for the settings to take effect
9. Validate the Java Agent is reporting to the controller by logging into the controller UI

For detailed information, see the link below:

[Install the AppDynamics Java Agent](#)

ABAP Agent Installation

The ABAP Agent needs to be installed on the SAP servers utilizing the ABAP stack.

There are four primary steps to perform, each with secondary steps involved. The four primary steps are:

1. Copy and unzip the ABAP Agent
2. Import the ABAP Agent Transports
3. Configure ABAP Agent and Install HTTP SDK
4. Activate Datavard Insight Collectors

Copy and Unzip ABAP Agent

The high-level steps to copy and unzip the ABAP Agent are listed below:

1. Ensure you are logged into the host server as the appropriate <SID>adm OS user
2. Copy the agent binary to a temporary directory on the server
3. Unzip the file into a temporary directory
4. Change the permissions on the contents of the agents' installation directory to give full read/write privileges at the owner and group levels

Import the ABAP Agent Transports



There are different versions of data files and cofiles within the ABAP agents' unzipped directory structure. The specific location of the appropriate files in the agents' directory structure to use will depend on the version of NetWeaver in use. For more information please see the link to the related documentation below:

[Install on SAP NetWeaver Systems](#)

The high-level steps to import the ABAP agent transports are listed below:

1. The ABAP Agents data files and cofiles should be copied from the temporary directories where they were unzipped over to the appropriate transport directories for the SAP module in question.
 - For example, for ECC we would copy the transports to " /usr/sap/trans/ECC2/cofiles" and " /usr/sap/trans/ECC2/data" , respectively
2. Set the permissions on the cofiles and data files to allow read/write access at the owner and group levels.
3. Log into the SAP system, execute transaction STMS:
 - Go to the import queue of the system where you want to install the ABAP agent
 - Select " Extras > Other Requests > Add" from the menu bar and add the vendor transports to the import queue one at a time
4. Import the transports in the appropriate order
 - The import order of the transports is specified in the "readme.txt" file of the ABAP Agent subdirectory that is specific to the version of NetWeaver in use

For more information please see the link to the related documentation below:

[Install on SAP NetWeaver Systems](#)

- Make sure that when selecting the " Execution" tab in the " Import Transport Request" pop-up dialog box to select the option " Synchronous" . When selecting the " Options" tab, put a checkmark next to " Ignore Invalid Component Version" .

Configure ABAP Agent / Install HTTP SDK



The steps below assume that your Linux hosts have glibc 2.5+ installed to allow for the automatic HTTP SDK installation. For more information please see the links to the related documentation below:

[Supported SAP Operating Systems](#)

[Installing HTTP SDK Automatically](#)

The high-level steps to configure the ABAP agent and install the HTTP SDK are listed below:

1. Log into the SAP system and execute transaction " /DVD/APPD_CUST" .
2. Switch to edit mode.
3. Fill in the fields on the screen to configure the agent to connect to the controller, SDK settings, and Analytics API settings..
4. Click the " Activate integration" button on the toolbar.
5. Click the " SDK Installation" button on the toolbar. This will take you to the " AppDynamics HTTP SDK Installation Manager" screen.
6. Select " Edit > Change Directory" from the menu bar.
 - a. Enter the path that was used for the agents' permanent base install directory (i.e. */usr/sap/appdyn*) in the field displayed in the pop-up dialog box shown below, and then click OK.
 - b. Click the " Install SDK" button on the toolbar.
 - c. Click the green checkmark to exit the screen and return to the " AppDynamics settings" screen.
 - d. Click the " Status" button on the toolbar. This will take you to the " AppDynamics status check" screen.
 - e. Click each " Start" button to start the HTTP SDK proxy on each SAP app server.

Activate Datavard Insight Collectors

Datavard Insights collect detailed performance data for an SAP system. It uses collector jobs that run as periodic background processes in the SAP system. These jobs must be scheduled to run.

Please refer to the links to the related documentation below:

[Datavard Insights Integration](#)

[Performance Collector Settings](#)

[SAP Dashboards](#)

[Mapping Between AppDynamics Metrics and Datavard Insights KPIs](#)

Server Visibility Agent Installation

The Server Visibility Agent must be installed on every application server and central services server that will be monitored.

The high-level steps for installing the Server Visibility agent are listed below:

1. Ensure you are logged into the host server as the appropriate <SID>adm OS user.
2. Create the permanent installation directory for the agent.

3. Unzip the file in the permanent directory (i.e. */usr/sap/appdyn/machine*).
4. Change the permissions on the contents of the agents' installation directory to give full read/write privileges at the owner and group levels.
5. Edit the configuration files in the agent installation directory to configure the agent to connect to the controller, provide the identity of the host, and so on.
6. Start the server visibility agent with the script provided in the agents' bin directory.
7. Validate the server visibility is reporting to the controller by logging into the controller UI.

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 15 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 20 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.
- Sarath Gonuganta, Cisco Systems Inc.