



Cisco UCS Integrated Infrastructure for SAP HANA

Design and Deployment of Cisco UCS Server and MapR Converged Data Platform with RedHat Enterprise Linux

Last Updated: April 30, 2018



About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, visit:

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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)

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	8
Solution Overview	9
Introduction	9
Audience	9
Purpose of this Document	9
What's New?	10
Solution Summary	10
Technology Overview	14
Cisco Unified Computing System	14
Cisco UCS Manager	14
Cisco UCS Fabric Interconnect	15
Cisco UCS 6332UP Fabric Interconnect	15
Cisco UCS 2304XP Fabric Extender	16
Cisco UCS Blade Chassis	16
Cisco UCS B480 M5 Blade Server	17
Cisco UCS C480 M5 Rack Servers	18
Cisco UCS C240 M5 Rack Servers	19
Cisco I/O Adapters for Blade and Rack-Mount Servers	20
Cisco VIC Interface Card	20
Cisco VIC 1380 Virtual Interface Card	21
Cisco VIC 1385 Virtual Interface Card	21
Cisco Unified Computing System Performance Manager	22
Cisco UCS Differentiators	23
MapR Converged Data Platform	24
Solution Design	27
SAP HANA System	27
Hardware Requirements for the SAP HANA Database	27
CPU	27
Memory	27
Network	28
Storage	30
Filesystem Layout	31
Operating System	32

High Availability	32
Physical Topology.....	33
Considerations.....	34
Scale	34
Performance.....	34
Deployment Hardware and Software	35
Configuration Guidelines.....	35
Topology	36
Physical Device Cabling.....	39
Software Revisions	46
Network Configuration	48
Cisco Nexus 3500 Series Switch Network Configuration	48
Cisco Nexus 3524 Initial Configuration	48
Enable Appropriate Cisco Nexus 3524 Switch Features and Settings	49
Create Global Policy to Enable Jumbo Frame and Apply the Policy to System Wide	50
Create VLANs for SAP HANA Traffic.....	50
Configure Network Interfaces Connecting to Cisco UCS Fabric Interconnect.....	50
Cisco Nexus 9000 Series Switch Network Configuration	51
Cisco Nexus 9000 A Initial Configuration	52
Cisco Nexus 9000 B Initial Configuration	53
Enable Appropriate Cisco Nexus 9000 Series Switches–Features and Settings	54
Create VLANs for SAP HANA Traffic.....	55
Configure Virtual Port-Channel Domain	55
Configure Network Interfaces for the VPC Peer Links	56
Configure Network Interfaces with Cisco UCS Fabric Interconnect	57
(Optional) Configure Network Interfaces for SAP HANA Backup/Data Source/Replication	59
(Optional) Management Plane Access for Cisco UCS Servers.....	60
Uplink into Existing Network Infrastructure	61
Cisco UCS Configuration	62
Initial Setup of Cisco UCS 6332 Fabric Interconnect.....	62
Cisco UCS 6332 Fabric Interconnect A.....	62
Cisco UCS 6332 Fabric Interconnect B.....	62
Cisco UCS for SAP HANA.....	63
Log in to Cisco UCS Manager	63
Upgrade Cisco UCS Manager Software to Version 3.2(2d)	63

Add Block of IP Addresses for KVM Access	63
Synchronize Cisco UCS to NTP.....	64
Cisco UCS Blade Chassis Connection Options	64
Enable Server and Uplink Ports.....	65
Configure Breakout Ports.....	66
Acknowledge Cisco UCS Chassis and Rack-Mount Servers	67
Create Uplink Port Channels	67
Create New Organization.....	71
Create MAC Address Pools	71
Create UUID Suffix Pool.....	73
Power Policy.....	74
Power Control Policy	74
Create Host Firmware Package	75
Create Local Disk Configuration Policy	76
Create Server BIOS Policy	79
Create Serial Over LAN Policy	84
Update Default Maintenance Policy	85
IPMI Access Profiles	85
Network Configuration	87
Set Jumbo Frames in Cisco UCS Fabric.....	87
Network Control Policy	88
Create Network Control Policy for Internal Network	89
LAN Configurations.....	90
Create VLANs.....	90
Create VLAN Groups	97
Create vNIC Template	103
Create a vNIC Template for Storage Network	106
Create a vNIC Template for AppServer Network.....	106
Create a vNIC Template for Backup Network.....	107
Create a vNIC Template for Client Network	107
Create a vNIC Template for DataSource Network	108
Create a vNIC Template for Replication Network	109
Create a vNIC Template for Management Network.....	109
Create a vNIC Template for IPMI Inband Access Network.....	110
Create a vNIC Template for MapR-01 Internal Network	111

Create a vNIC template for MapR-02 Internal Network.....	111
Create Boot Policies	112
Create IP Pool for Inband Management.....	113
Configure the Inband Profile	114
Create Service Profile Templates for SAP HANA Scale-Out Servers.....	115
Create Service Profile from the Template	125
Create Service Profile Templates for MapR Servers.....	127
Create Service Profiles	133
MapR Storage Configuration	135
MapR Server RAID Configuration	135
MapR Server Operating System Installation	143
Operating System Network Configuration	153
Configure the Network.....	154
DNS.....	156
RedHat System Update and OS Customization for MapR Servers	157
Install Cisco eNIC Driver	158
Network Time	159
SSH Keys.....	160
MapR Installation	160
Virtual IP Range	161
Preparing Online Repository	161
Installing MapR Packages	161
Initial Configuration	161
Starting up MapR cluster	163
Add License to MapR Cluster	163
Create Virtual IPs	164
Create Volumes	165
HANA System Configuration	168
RHEL Operating System Installation	168
Operating System Network Configuration	177
DNS.....	181
Hosts file	181
RedHat System Update and OS customization for SAP HANA	183
Install Cisco eNIC Driver	186
Network Time	187

SSH Keys.....	187
Mount Options.....	187
SAP HANA Installation.....	189
Important SAP Notes.....	189
SAP HANA IMDB Related Notes.....	189
Linux Related Notes.....	190
SAP Application Related Notes.....	190
Third Party Software.....	191
SAP HANA Virtualization.....	191
High Availability (HA) Configuration for Scale-Out.....	191
High Availability Configuration.....	191
Enable the SAP HANA Storage Connector API.....	193
Test the IPMI Connectivity.....	194
SAP HANA Parameters for Scale-Out.....	194
SAP HANA Hardware Configuration Check Tool Parameter.....	195
About the Authors.....	197
Acknowledgements.....	197

Executive Summary

Organizations in every industry are generating and using more data than ever before; from customer transactions and supplier delivery considerations, to real-time user-consumption statistics. Without scalable infrastructure that can store, process, and analyze big data sets in real time, companies are unable to use this information to their advantage. The Cisco UCS Integrated Infrastructure for SAP HANA Scale-Out with **the Cisco Unified Computing System™ (Cisco UCS) helps companies easily harness information and make** better business decisions that let them stay ahead of the competition. Our solutions help improve access to all of your data, accelerate business decision making with policy-based, simplified management, lower deployment risk, and reduce total cost of ownership (TCO). Our innovations give you the key to unlock the intelligence in your data and interpret it with a new dimension of context and insight to help you create a sustainable, competitive business advantage.

Cisco Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. The Cisco UCS Integrated Infrastructure for SAP HANA with MapR Converged Data Platform provides an end-to-end architecture that demonstrate support for multiple SAP HANA workloads with high availability and server redundancy.

The solution consists of Cisco UCS B-series B480-M5 blade servers and or Cisco UCS C-series C480-M5 rack mount servers as compute nodes for SAP HANA and Cisco UCS C-Series C240 rack mount servers as storage nodes. The next generation Cisco UCS Fabric Interconnect 6332 with 40 Gb Ethernet for Server Management and Storage Connectivity. Cisco UCS service profiles enable rapid and consistent server configuration, and automation simplifies ongoing system maintenance activities such as deploying firmware updates across the entire cluster as a single operation. Advanced monitoring capabilities raise alarms and send notifications about the health of the entire cluster so that you can proactively address concerns before they affect data analysis. The Storage Nodes are composed of Cisco UCS Servers with MapR Converged Data Platform, which is a modern NFS-mountable distributed filesystem. MapR-FS is a complete POSIX file system that handles raw disk I/O for big data workload with direct access to storage hardware which dramatically improving performance and scale to thousands of nodes, and trillions of files, with extremely high throughput. MapR-FS includes enterprise-grade features such as block-level mirroring for mission-critical disaster recovery as well as load balancing, and consistent snapshots for easy data recovery.

Solution Overview

Introduction

Cisco UCS Integrated Infrastructure provides a pre-validated, ready-to-deploy infrastructure, which reduces the time and complexity involved in configuring and validating a traditional data center deployment. Cisco UCS Platforms is flexible, reliable and cost effective to facilitate various deployment options of the applications while being easily scalable and manageable. The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a SAP HANA solution. This document describing the infrastructure installation and configuration to run SAP HANA on a dedicated or shared infrastructure.

SAP HANA is SAP SE's implementation of in-memory database technology. The SAP HANA database takes advantage of the low cost main memory (RAM), data-processing capabilities of multicore processors, and faster data access to provide better performance for analytical and transactional applications. SAP HANA offers a multi-engine, query-processing environment that supports relational data (with both row- and column-oriented physical representations in a hybrid engine) as well as a graph and text processing for semi-structured and unstructured data management within the same system. As an appliance, SAP HANA combines software components from SAP optimized for certified hardware. However, this solution has a preconfigured hardware set-up and preinstalled software package that is dedicated for SAP HANA. In 2013, SAP introduced SAP HANA Tailored Datacenter Integration (TDI) option; TDI solution offers a more open and flexible way for integrating SAP HANA into the data center by reusing existing enterprise storage hardware, thereby reducing hardware costs. With the introduction of SAP HANA TDI for shared infrastructure, the Cisco UCS Integrated Infrastructure solution provides the advantage of having the compute, storage, and network stack integrated with the programmability of the Cisco Unified Computing System (Cisco UCS). SAP HANA TDI option enables organizations to run multiple SAP HANA production systems on a shared infrastructure. It also enables customers to run the SAP applications servers and SAP HANA database hosted on the same infrastructure.

For more information about SAP HANA, see the SAP Help Portal: <http://help.sap.com/hana/>

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers deploying the Cisco Integrated Infrastructure for SAP HANA with MapR Converged Data Platform. External references are provided wherever applicable, but readers are expected to be familiar with the technology, infrastructure, and database security policies of the customer installation.

Purpose of this Document

This document describes the steps required to deploy and configure a Cisco Datacenter Solution for SAP HANA. **Cisco's validation provides further confirmation with regard to component compatibility**, connectivity and correct operation of the entire integrated stack. This document showcases one of the variants of Cisco Integrated Infrastructure for SAP HANA. While readers of this document are expected to have sufficient

knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are provided in this CVD.

What's New?

Cisco UCS Integrated Infrastructure for SAP HANA solution designed with next generation Fabric Interconnect which provides 40GbE ports. The solution is designed with 40GbE end-to-end network including Storage network. The persistent storage is configured on Cisco UCS C240 C-series servers with MapR Converged Data Platform. MapR-FS provides distributed, reliable, high performance, scalable, and full read/write data storage for SAP HANA.

Solution Summary

The Cisco UCS Integrated Infrastructure for SAP HANA with MapR Converged Data Platform provides an end-to-end architecture with Cisco Hardware that demonstrate support for multiple SAP HANA workloads with high availability and server redundancy. The solution supports up to 16 x Cisco UCS B480-M5 B-series blade servers or 16 x Cisco UCS C480-M5 C-series rack mount servers for SAP HANA and up to 8 x Cisco UCS C240-M5 C-Series rack mount servers for storage. The next Gen Cisco UCS Fabric Interconnect 6332 with 40 GbE network bandwidth for Server Management and Network connectivity. The Cisco UCS C240-M5 servers provides persistent storage with MapR Converged Data Platform, which is a modern NFS-mountable distributed file-system. The Nexus 3000 series Ethernet switch is used for failover purpose. In case of 2304 Fabric Expander failure or link failure between Server and Fabric Interconnect, the data traffic path will use Nexus 3000 series switch for High Availability and redundancy. Figure 1 shows the Cisco UCS Integrated Infrastructure for SAP HANA block diagram with Cisco UCS C480-M5.



The reference architecture documented in this CVD consists of 8 x Cisco UCS B/C480-M5 servers for SAP HANA and 4 x Cisco UCS C240-M5 C-Series rack mount servers for storage.

Figure 1 Cisco UCS Integrated Infrastructure for SAP HANA with Cisco UCS C480 M5






Figure 2 shows the Cisco UCS Integrated Infrastructure for SAP HANA block diagram with Cisco UCS B480 M5.

Figure 2 Cisco UCS Integrated Infrastructure for SAP HANA with Cisco UCS B480 M5



The solution can be designed with a pair for Cisco Nexus 9000 series switches, alternative to single Nexus 3000 series switch. Two Cisco Nexus 9000 series switches is configured with vPC between network switch and Cisco Fabric Interconnect as show in Figure 3.

Figure 3 Cisco UCS B/C480 M5 Scale-Out for SAP HANA Pair of Nexus Switches



Technology Overview

Cisco Unified Computing System

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

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

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

Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management of all software and hardware components of **the Cisco Unified Computing System™ (Cisco UCS) and Cisco HyperFlex™ Systems across multiple chassis** and rack servers and thousands of virtual machines. It supports all Cisco UCS product models, including Cisco UCS B-Series Blade Servers and C-Series Rack Servers, Cisco UCS Mini, and Cisco HyperFlex hyperconverged infrastructure, as well as the associated storage resources and networks. Cisco UCS Manager is embedded on a pair of Cisco UCS 6300 or 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The manager participates in server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection.

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

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

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

Cisco UCS Fabric Interconnect

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

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

Cisco UCS 6332UP Fabric Interconnect

The Cisco UCS 6332 Fabric Interconnect is the management and communication backbone for Cisco UCS B-Series Blade Servers, C-Series Rack Servers, and 5100 Series Blade Server Chassis. All servers attached to 6332 Fabric Interconnects become part of one highly available management domain. The Cisco UCS 6332UP 32-Port Fabric Interconnect is a 1-rack-unit 40 Gigabit Ethernet, FCoE and Fibre Channel switch offering up to 2.56 Tbps throughput and up to 32 ports. The switch has 32 fixed 40-Gbps Ethernet and FCoE

ports. Cisco UCS 6332UP 32-Port Fabric Interconnect have ports that can be configured for the breakout feature that supports connectivity between 40 Gigabit Ethernet ports and 10 Gigabit Ethernet ports. This feature provides backward compatibility to existing hardware that supports 10 Gigabit Ethernet. A 40 Gigabit Ethernet port can be used as four 10 Gigabit Ethernet ports. Using a 40 Gigabit Ethernet SFP, these ports on a Cisco UCS 6300 Series Fabric Interconnect can connect to another fabric interconnect that has four 10 Gigabit Ethernet SFPs.

Figure 4 Cisco UCS 6332 UP Fabric Interconnect



Cisco UCS 2304XP Fabric Extender

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

Figure 5 Cisco UCS 2304 XP



Cisco UCS Blade Chassis

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

The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors. Four hot-swappable power supplies are accessible from the front of the chassis, and single-phase AC, -48V DC, and 200 to 380V DC power supplies and chassis are available. These power supplies are up to 94 percent efficient and meet the requirements for the 80 Plus Platinum rating. The power subsystem can be configured to support nonredundant, N+1 redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays that can support either Cisco UCS 2000 Series Fabric Extenders or the Cisco UCS 6324 Fabric Interconnect. A passive midplane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots.

The Cisco UCS Blade Server Chassis is shown in Figure 6.

Figure 6 Cisco Blade Server Chassis (front and back view)



Cisco UCS B480 M5 Blade Server

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

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

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

The Cisco UCS B480 M5 Blade Server offers:

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

Figure 7 Cisco UCS B480 M5 Blade Server



Cisco UCS C480 M5 Rack Servers

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

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

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

Figure 8 Cisco UCS C480 M5 Rack Server



Cisco UCS C240 M5 Rack Servers

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

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

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

Cisco UCS C240 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. When used in combination with Cisco UCS Manager, the Cisco UCS C240 M5 brings the power and automation of unified computing to enterprise applications, including Cisco® SingleConnect technology, drastically reducing switching and cabling requirements. Cisco UCS Manager uses service profiles, templates, and policy-based management to enable rapid deployment and help ensure deployment consistency. It also enables end-to-end server visibility, management, and control in both virtualized and bare-metal environments.

Figure 9 Cisco UCS C240 M5 Rack Server



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


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

Cisco VIC Interface Card

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

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

Figure 10 Cisco UCS 1340 VIC Card



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

Cisco VIC 1380 Virtual Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1380 is a dual-port 40-Gbps Ethernet, or dual 4 x 10 Fibre Channel over Ethernet (FCoE)-capable mezzanine card designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers. The card enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1380 supports Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.


Figure 11 Cisco UCS 1380 VIC Card



Cisco VIC 1385 Virtual Interface Card

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

Figure 12 Cisco UCS 1385 VIC Card



Cisco Unified Computing System Performance Manager

Cisco UCS Performance Manager is a purpose-built data center operations management solution. It unifies the monitoring of key applications, business services, and integrated infrastructures across dynamic, heterogeneous, physical, and virtual Cisco UCS-powered data centers. Cisco UCS Performance Manager uses Cisco UCS APIs to collect data from Cisco UCS Manager to display comprehensive information about all Cisco UCS infrastructure components. With a customizable view, data center staff can see application services and view performance and component or service availability information for Cisco UCS integrated infrastructures.

Cisco UCS Performance Manager dynamically collects information about Cisco UCS servers, network, storage, and virtual machine hosts using an agentless information gathering approach. The solution provided the following:

- Unifies performance monitoring and management of Cisco UCS integrated infrastructure solutions
- Delivers real-time views of fabric and data center switch bandwidth usage and capacity thresholds
- Discovers and creates a relationship model of each system, giving staff a single, accurate view of all components
- Allows staff to navigate into individual Cisco UCS infrastructure components when troubleshooting and resolving issues

Cisco UCS Performance Manager provides deep visibility of Cisco UCS integrated infrastructure performance for service profiles, chassis, fabric extenders, adapters, virtual interface cards, ports, and uplinks for granular data center monitoring. Customers can use Cisco UCS Performance Manager to maintain service-level agreements (SLAs) by managing optimal resource allocation to prevent under-provisioning and avoid performance degradation. By defining component or application-centric views of critical resources, administrators can monitor SLA health and performance from a single console, eliminating the need for multiple tools.

For detailed information, see the [Cisco UCS Performance Manager Install Guide](#).

Cisco UCS Differentiators

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


- **Embedded management:** In Cisco Unified Computing System, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers. Also, a pair of FIs can manage up to 40 chassis, each containing 8 blade servers. This gives enormous scaling on management plane.
- **Unified fabric:** In Cisco Unified Computing System, from blade server chassis or rack server fabric extender to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O, results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.
- **Auto discovery:** By simply inserting the blade server in the chassis or connecting rack server to the fabric extender, discovery and inventory of compute resource occurs automatically without any management intervention. Combination of unified fabric and auto-discovery enables wire-once architecture of Cisco Unified Computing System, where compute capability of Cisco Unified Computing System can extend easily, while keeping the existing external connectivity to LAN, SAN and management networks.
- **Policy based resource classification:** When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD focuses on the policy-based resource classification of Cisco UCS Manager.
- **Combined Rack and Blade server management:** Cisco UCS Manager can manage Cisco UCS B-Series Blade Servers and Cisco UCS C-Series Rack Servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic. This CVD focuses on the combination of B-Series and C-Series Servers to demonstrate stateless and form factor independent computing work load.
- **Model-based management architecture:** Cisco UCS Manager Architecture and management database is model based and data driven. Open, standard based XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management system, such as VMware vCloud director, Microsoft system center, and Citrix CloudPlatform.
- **Policies, Pools, Templates:** Management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- **Loose referential integrity:** In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each-other. This provides great flexibilities where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.

- Policy resolution: In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to other policy by name is resolved in the org hierarchy with closest policy match. If no policy with specific name is found in the hierarchy till root org, **then special policy named “default” is searched. This policy** resolution practice enables automation friendly management APIs and provides great flexibilities to owners of different orgs.
- Service profiles and stateless computing: Service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in multi-tenancy support: Combination of policies, pools and templates, loose referential integrity, policy resolution in org hierarchy and service profile based approach to compute resources make Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- Virtualization aware network: VM-FEX technology makes access layer of network aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-**profiles defined by the network administrators’ team.** VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
- Simplified QoS: Even though fibre-channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

MapR Converged Data Platform

The MapR Converged Data Platform solves the crisis of complexity that results from continually deploying workload-specific data silos. Within a single platform on a single codebase, it converges the key technologies that make up a modern data architecture, including a distributed file system, a multi-model NoSQL database, a publish/subscribe event streaming engine, ANSI SQL, and a broad set of open source data management and analytics technologies.

The MapR Converged Data Platform delivers speed, scale, and reliability, driving both operational and analytical workloads in a single platform. It is architected with many performance optimizations to get the **most out of your hardware. It efficiently scales horizontally (“scale out”) on commodity hardware to cost-effectively** expand or contract your computing power as your load changes, even to exabyte levels. It provides mission-critical high availability (HA), disaster recovery (DR), and data recovery features to maximize uptime and reduce risk of data loss.



The MapR Converged Data Platform includes the following components:

- **Web-Scale Storage.** MapR-FS is a distributed POSIX file system with full read-write semantics, which can scale to exabytes of data and trillions of files in a single cluster.
- **NoSQL Database.** MapR-DB is a multi-model NoSQL database that natively supports JSON document and wide column data models with high performance, consistent low latency, strong consistency, multi-master replication, granular security, and completely automatic self-tuning.
- **Event Streaming.** MapR Streams is a publish-subscribe, event stream transport engine for reliably delivering ordered messages at high volumes and velocities.
- **Apache™ Hadoop®.** MapR provides open source ecosystem projects to handle a variety of big data management tasks. Projects include Apache Storm™, Apache Pig, Apache Hive™, Apache Mahout™, YARN, Apache Sqoop™, Apache Flume™, and more.
- **Apache Spark™.** MapR provides the full stack of the popular Spark tool set for fast, in-memory processing of big data.
- **ANSI SQL.** Apache Drill™ is a SQL query engine that provides low-latency results, using familiar **business intelligence (BI) tools**. It also queries “**schemaless**” data such as JSON to enable self-service data exploration and analytics.
- **Third-party compute engines and custom apps.** Due to interoperability features built into the system, a wide variety of third-party compute engines can run on MapR to take advantage of its speed, scale, and reliability at the data storage level.

The Cisco UCS C240 based Storage for SAP HANA takes advantage of the enterprise grade MapR-FS in MapR Platform Services. It offers the appliance a robust storage layer that is highly available, resilient and performant.

MapR converged data platform is backed by the robust MapR-FS, which can be accessed through the NFS gateway. The SAP HANA Servers mounts MapR-FS using NFS client. Data can be persisted to the storage

system managed by MapR-FS. MapR-FS is distributed, has a global name space, real-time read/write access, volume based, secure and has many other benefits compared to HDFS.

For additional information, please see the [MapR File System](#).

Solution Design

This section describes the SAP HANA system requirements defined by SAP and Architecture of Cisco UCS Integrated Infrastructure for SAP HANA.

SAP HANA System

SAP HANA System on a Single Server Scale-Up, is the simplest of the installation types. It is possible to run an SAP HANA system entirely on one host and then scale the system up as needed. All data and processes are located on the same server and can be accessed locally. The network requirements for this option minimum one 1-Gb Ethernet (access) and one 10-Gb Ethernet storage networks are sufficient to run SAP HANA scale-up. SAP HANA Scale-Out option is used if the SAP HANA system does not fit into the main memory of a single server based on the rules defined by SAP. In this method, multiple independent servers are combined to form one system and the load is distributed among multiple servers. In a distributed system, each index server is usually assigned to its own host to achieve maximum performance. It is possible to assign different tables to different hosts (partitioning the database), or a single table can be split across hosts (partitioning of tables). SAP HANA Scale-Out supports failover scenarios and high availability. Individual hosts in a distributed system have different roles master, worker, slave, standby depending on the task.



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

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

Hardware Requirements for the SAP HANA Database

There are hardware and software requirements defined by SAP to run SAP HANA systems. This Cisco Validated Design uses guidelines provided by SAP.

For additional information, go to: <http://saphana.com>.



This document does not cover the updated information published by SAP.

CPU

SAP HANA supports servers equipped with Intel(R) Xeon(R) Platinum 8176, 8176M, 8180 and 8180M CPU.

Memory

SAP HANA Scale-Out solution is supported in the following memory configurations:

- Homogenous symmetric assembly of dual in-line memory modules (DIMMs) for example, DIMM size or speed should not be mixed
- Maximum use of all available memory channels
- Memory of 1.5 TB or 3 TB per 4 Socket Server for SAP NetWeaver Business Warehouse (BW) and DataMart

Network

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

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

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

Table 1 lists all the networks defined by SAP or Cisco or requested by customers.

Table 1 List of Known Networks


Name	Use Case	Solutions	Bandwidth Requirements	Solution Design
Client Zone Networks				
Application Server Network	SAP Application Server to DB communication	All	1 or 10 GbE	10 or 40 GbE
Client Network	User / Client Application to DB communication	All	1 or 10 GbE	10 or 40 GbE
Data Source Network	Data import and external data integration	Optional for all SAP HANA systems	1 or 10 GbE	10 or 40 GbE
Internal Zone Networks				
Inter-Node Network	Node to node communication within a scale-out configuration	Scale-Out	10 GbE	40 GbE

Name	Use Case	Solutions	Bandwidth Requirements	Solution Design
System Replication Network	SAP HANA System Replication	For SAP HANA Disaster Tolerance	TBD with Customer	TBD with Customer
Storage Zone Networks				
Backup Network	Data Backup	Optional for all SAP HANA systems	10 GbE	10 or 40 GbE
Storage Network	Node to Storage communication	All	10 GbE	20 or 40 GbE
Infrastructure Related Networks				
Administration Network	Infrastructure and SAP HANA administration	Optional for all SAP HANA systems	1 GbE	10 or 40 GbE
Boot Network	Boot the Operating Systems through PXE/NFS or FCoE	Optional for all SAP HANA systems	1 GbE	N/A

For detailed information about the network requirements for SAP HANA see: [SAP HANA Network Requirements](#).

The network need to be properly segmented and must be connected to the same core/ backbone switch as shown in Figure 13 **based on customer's high**-availability and redundancy requirements for different SAP HANA network segments.

Figure 13 High-Level SAP HANA Network Overview



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

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



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

Storage

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

<http://scn.sap.com/docs/DOC-48516>.

All relevant information about storage requirements is documented in the white paper [SAP HANA Storage Requirements](#).

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


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

Filesystem Layout

Figure 14 shows the file system layout and the required storage sizes to install and operate SAP HANA. For the Linux OS installation (/root) 10 GB of disk size is recommended. Additionally, 50 GB must be provided for the /usr/sap since the volume used for SAP software that supports SAP HANA.

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

Figure 14 File System Layout for 2 Node Scale-Out System



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

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

Root-FS: 10 GB

/usr/sap: 50 GB

The installation binaries, trace and configuration files are stored on a shared filesystem, which should be accessible for all hosts in the distributed installation. The size of shared filesystem should be equal to one times memory in each host for every four worker nodes. For example: In a distributed installation with eight hosts with 3 TB of memory each, shared file system should be 6 TB.

For each HANA host there should be a mount point for data and log volume.

Size of the file system for data volume for Appliance option is three times the host memory:

/hana/data/<sid>/mntXXXXX: 3 x Memory

Size of the file system for data volume with TDI option is one times the host memory:

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

The Log volume must be as follows:

- **Half of the server memory for systems \leq 256 GB memory**
- **Min 512 GB for systems with \geq 512 GB memory**

Operating System

The supported operating systems for SAP HANA are as follows:

- SUSE Linux Enterprise Server for SAP Applications
- RedHat Enterprise Linux for SAP HANA



This document provides the installation process for the RedHat Enterprise Linux for SAP Applications only. For SUSE Linux Enterprise Server option please follow the CVD [UCS II for SAP HANA with SLES](#).

High Availability

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

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

SAP HANA Scale-Out comes with an integrated high-availability function. If a SAP HANA system is configured with a stand-by node, a failed part of SAP HANA will start on the stand-by node automatically.

For automatic host failover, storage connector API must be properly configured for the implementation and operation of the SAP HANA.

Please check the latest information from SAP at: <http://saphana.com> or <http://service.sap.com/notes>.

Physical Topology

The Cisco UCS Integrated Infrastructure for SAP HANA with MapR Converged Data Platform provides an end-to-end architecture with Cisco Hardware that demonstrate support for multiple SAP HANA workloads with high availability and server redundancy. The architecture uses Cisco UCS Manager with combined Cisco UCS B-Series and C-Series Servers with Cisco UCS Fabric Interconnect. The uplink from Cisco UCS Fabric Interconnect is connected to Nexus 3524 switches for High Availability and Failover functionality. The Cisco UCS C-Series Rack Servers are connected directly to Cisco UCS Fabric Interconnect with single-wire management feature, the data traffic between HANA servers and Storage will be contained in the Cisco UCS Fabric Interconnect. This infrastructure is deployed to provide IP based Storage access using NFS protocols with file-level access to shared storage.

The Base Reference architecture includes:

- Cisco Unified Computing System
 - 2 x Cisco UCS 6332UP 32-Port Fabric Interconnects
 - 2 x Cisco UCS 5108 Blade Chassis with 2 x Cisco UCS 2304 Fabric Extenders with 4x 40 Gigabit Ethernet interfaces
 - 8 x Cisco UCS B480 M5 High-Performance Blade Servers with 1x Cisco UCS Virtual Interface Card (VIC) 1380 and 1x Cisco UCS Virtual Interface Card (VIC) 1340

Or

 - 8 x Cisco UCS C480 M5 High-Performance Rack Servers with 1x Cisco UCS Virtual Interface Card (VIC) 1385
 - 4 x Cisco UCS C240 M5 High-Performance Blade Servers with Cisco UCS Virtual Interface Card (VIC) 1385
- Cisco Network
 - For Base design:
 - 1 x Cisco Nexus 3524 Switch for 10 Gigabit Ethernet connectivity between the two Cisco UCS Fabric Interconnects for failover scenarios.

The solution can be designed for Data Center Design option using a pair for Cisco Nexus 9000 series switches. With two Cisco Nexus switches, vPC is configured between Cisco Nexus 9000 series switches and Cisco Fabric Interconnect.

- For Enterprise Data Center Design option:
 - 2 x Cisco Nexus 9372 Switch for 10 Gigabit Ethernet connectivity between the two Cisco UCS Fabric Interconnects

Considerations

Scale

Although this is the base design, each of the components can be scaled easily to support specific business requirements. Additional servers or even blade chassis can be deployed to increase compute capacity without additional Network components. Two Cisco UCS 6332UP, 32 port Fabric interconnect can support up to:

- 16 Cisco UCS B-Series B480 M5 Server with 8 Blade Server Chassis

Or

- 16 Cisco UCS C-Series C480 M5 Server
- 12 Cisco UCS C240 M5 Server

A minimum of 4 x Cisco UCS C240 M5 Server is required install MapR Converged Data Platform with High Availability. 4 x Cisco UCS C240 M5 can support up to 14 Active HANA Servers. The scaling of storage to computer node is linear, for every 7 x Cisco UCS Server for Active HANA Servers, 2 x Cisco UCS C240 M5 for Storage is required to meet the SAP HANA TDI storage performance defined by SAP SE.

Performance

The solution is designed to meet SAP HANA performance requirement defined by SAP SE. The HANA Server and Storage server are connecting to Cisco UCS Fabric Interconnect, all the data traffic between HANA node and Storage node is contained in the Cisco UCS Fabric Interconnect. Each HANA Server and Storage Server are equipped with 2 x 40GbE capable Cisco Virtual Interface Cards, the storage network provides dedicated bandwidth between HANA servers and Storage Servers. The traffic shaping and Quality of Service (QoS) is configured and managed by Cisco UCS Fabric Interconnect. For HANA node-to-node network, 40 Gb dedicated network bandwidth is provided with non-blocking mode. The Cisco UCS Integrated Infrastructure for SAP HANA quarantines the bandwidth and latency for best performance to run SAP HANA.

Deployment Hardware and Software

Configuration Guidelines

This document provides the details to configure a fully redundant, highly available configuration for a Cisco UCS B/C480 M5 Scale-Out for SAP HANA.

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

Table 2 Configuration Variables

Variable	Description	Customer Implementation Value
<<var_nexus_HA_hostname>>	Cisco Nexus 3524 HA Switch host name	
<<var_nexus_HA_mgmt0_ip>>	Out-of-band Cisco Nexus 3524 HA Switch management IP address	
<<var_nexus_HA_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_HA_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_oob_vlan_id>>	Out-of-band management network VLAN ID	
<<var_mgmt_vlan_id>>	Management network VLAN ID	
<<var_nexus_vpc_domain_mgmt_id>>	Unique Cisco Nexus switch VPC domain ID for Management Switch	
<<var_nexus_vpc_domain_id>>	Unique Cisco Nexus switch VPC domain ID	
<<var_nexus_A_hostname>>	Cisco Nexus 9000 A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus 9000 A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus 9000 B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus 9000 B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_mapr-01_vlan_id>>	MapR Internal network 01 VLAN ID	

Variable	Description	Customer Implementation Value
<<var_mapr-02_vlan_id>>	MapR Internal network 02 VLAN ID	
<<var_storage_vlan_id>>	Storage network for HANA Data/log VLAN ID	
<<var_internal_vlan_id>>	Node to Node Network for HANA Data/log VLAN ID	
<<var_backup_vlan_id>>	Backup Network for HANA Data/log VLAN ID	
<<var_client_vlan_id>>	Client Network for HANA Data/log VLAN ID	
<<var_appserver_vlan_id>>	Application Server Network for HANA Data/log VLAN ID	
<<var_datasource_vlan_id>>	Data source Network for HANA Data/log VLAN ID	
<<var_replication_vlan_id>>	Replication Network for HANA Data/log VLAN ID	
<<var_inband_vlan_id>>	In-band management network VLAN ID	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	

Topology

The information in this section is provided as a reference for cabling the network and compute components. For connectivity between 40GbE ports to 10GbE ports on Cisco Switches, Cisco QSFP to Four SFP+ Copper Breakout Cables are used. These breakout cables connect to a 40G QSFP port of a Cisco Fabric Interconnect on one end and to four 10G SFP+ ports of a Cisco switch on the other end

Figure 15 illustrates the cabling topology for Cisco UCS Integrated Infrastructure for SAP HANA configuration using the Cisco Nexus 3000.

Figure 15 Cabling Topology for Cisco UCS Integrated Infrastructure for SAP HANA




Figure 16 Cisco UCS Integrated Infrastructure for SAP HANA Configuration using the Cisco Nexus 3000





Figure 17 shows the cabling topology for Cisco UCS Integrated Infrastructure for SAP HANA configuration using the pair of Cisco Nexus 9000 series switches.

Figure 17 Cabling Topology for Cisco UCS Integrated Infrastructure for SAP HANA



Physical Device Cabling

The information in this section is provided as a reference for cabling the network and compute components.

To simplify cabling requirements, the tables include both local and remote device and port locations. The following tables show the out-of-band management ports connectivity into preexisting management

infrastructure, the Management Ports cabling needs to be adjusted accordingly. These Management interfaces will be used in various configuration steps.

Table 3 through 0provides the details of all the connections.

Table 3 Cisco UCS Fabric Interconnect A - Cabling Information for Cisco UCS B480-M5 Server(s) Option

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3/1	40GbE	Cisco Nexus 3000 HA	Eth 1/3
	Eth1/3/2	QSFP to 4 SFP+ break-out cables	Cisco Nexus 3000 HA	Eth 1/4
	Eth1/3/3		Cisco Nexus 3000 HA	Eth 1/5
	Eth1/3/4		Cisco Nexus 3000 HA	Eth 1/6
	Eth1/3/1*		40GbE	Cisco Nexus 9000 A
	Eth1/3/2*	QSFP to 4 SFP+ break-out cables	Cisco Nexus 9000 A	Eth 1/4
	Eth1/3/3*		Cisco Nexus 9000 B	Eth 1/3
	Eth1/3/4*		Cisco Nexus 9000 B	Eth 1/4
	Eth1/4		40GbE	
	Eth1/5	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM 1/1
	Eth1/6	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM 1/2
	Eth1/7	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM 1/3
	Eth1/8	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM 1/4
	Eth1/9	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/1
	Eth1/10	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/2
	Eth1/11	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/3
	Eth1/12	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/4
	Eth1/13	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/1
	Eth1/14	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/2
	Eth1/15	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/3
	Eth1/16	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/4
	Eth1/17	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) A	IOM 1/1
Eth1/18	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) A	IOM 1/2	
Eth1/19	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) A	IOM 1/3	
Eth1/20	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) A	IOM 1/4	

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/21	40GbE	Cisco UCS C240-M5-1	VIC 1385 Port 1
	Eth1/22	40GbE	Cisco UCS C240-M5-2	VIC 1385 Port 1
	Eth1/23	40GbE	Cisco UCS C240-M5-3	VIC 1385 Port 1
	Eth1/24	40GbE	Cisco UCS C240-M5-4	VIC 1385 Port 1
	Eth1/25	40GbE	Cisco UCS C240-M5-5	VIC 1385 Port 1
	Eth1/26	40GbE	Cisco UCS C240-M5-6	VIC 1385 Port 1
	Eth1/27	40GbE	Cisco UCS C240-M5-7	VIC 1385 Port 1
	Eth1/28	40GbE	Cisco UCS C240-M5-8	VIC 1385 Port 1
	MGMT0	GbE	Customer's Management Switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

* The ports ETH1/3/1-4 are used with pair Nexus 9000 Switches design option.

Table 4 Cisco UCS Fabric Interconnect B - Cabling Information for Cisco UCS B480 M5 Server(s) Option

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3/1	40GbE	Cisco Nexus 3000 HA	Eth 1/7
	Eth1/3/2	QSFP to 4 SFP+ break-out cables	Cisco Nexus 3000 HA	Eth 1/8
	Eth1/3/3		Cisco Nexus 3000 HA	Eth 1/9
	Eth1/3/4		Cisco Nexus 3000 HA	Eth 1/10
	Eth1/3/1*		40GbE	Cisco Nexus 9000 A
	Eth1/3/2*	QSFP to 4 SFP+ break-out cables	Cisco Nexus 9000 A	Eth 1/6
	Eth1/3/3*		Cisco Nexus 9000 B	Eth 1/5
	Eth1/3/4*		Cisco Nexus 9000 B	Eth 1/6
	Eth1/4		40GbE	
	Eth1/5	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/1
	Eth1/6	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/2
	Eth1/7	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/3
	Eth1/8	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/4
Eth1/9	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/1	

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/10	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/2
	Eth1/11	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/3
	Eth1/12	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/4
	Eth1/13	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/1
	Eth1/14	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/2
	Eth1/15	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/3
	Eth1/16	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/4
	Eth1/17	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) B	IOM 1/1
	Eth1/18	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) B	IOM 1/2
	Eth1/19	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) B	IOM 1/3
	Eth1/20	40GbE	Cisco UCS Chassis 4 Fabric Extender (FEX) B	IOM 1/4
	Eth1/21	40GbE	Cisco UCS C240-M5-1	VIC 1385 Port 2
	Eth1/22	40GbE	Cisco UCS C240-M5-2	VIC 1385 Port 2
	Eth1/23	40GbE	Cisco UCS C240-M5-3	VIC 1385 Port 2
	Eth1/24	40GbE	Cisco UCS C240-M5-4	VIC 1385 Port 2
	Eth1/25	40GbE	Cisco UCS C240-M5-5	VIC 1385 Port 2
	Eth1/26	40GbE	Cisco UCS C240-M5-6	VIC 1385 Port 2
	Eth1/27	40GbE	Cisco UCS C240-M5-7	VIC 1385 Port 2
	Eth1/28	40GbE	Cisco UCS C240-M5-8	VIC 1385 Port 2
	MGMT0	GbE	Customer's Management Switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1
	L2	GbE	Cisco UCS fabric interconnect A	L2

* The ports ETH1/3/1-4 are used with pair Nexus 9000 Switches design option.

Table 5 Cisco UCS Fabric Interconnect A - Cabling Information for Cisco UCS C480 M5 Server(s) Option

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3/1	40GbE QSFP to 4 SFP+ break- out cables	Cisco Nexus 3000 HA	Eth 1/3
	Eth1/3/2		Cisco Nexus 3000 HA	Eth 1/4
	Eth1/3/3		Cisco Nexus 3000 HA	Eth 1/5
	Eth1/3/4		Cisco Nexus 3000 HA	Eth 1/6

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/3/1*	40GbE QSFP to 4 SFP+ break- out cables	Cisco Nexus 9000 A	Eth 1/3
	Eth1/3/2*		Cisco Nexus 9000 A	Eth 1/4
	Eth1/3/3*		Cisco Nexus 9000 B	Eth 1/3
	Eth1/3/4*		Cisco Nexus 9000 B	Eth 1/4
	Eth1/4	40GbE		
	Eth1/5	40GbE	Cisco UCS C480-M5-1	VIC 1385 Port 1
	Eth1/6	40GbE	Cisco UCS C480-M5-2	VIC 1385 Port 1
	Eth1/7	40GbE	Cisco UCS C480-M5-3	VIC 1385 Port 1
	Eth1/8	40GbE	Cisco UCS C480-M5-4	VIC 1385 Port 1
	Eth1/9	40GbE	Cisco UCS C480-M5-5	VIC 1385 Port 1
	Eth1/10	40GbE	Cisco UCS C480-M5-6	VIC 1385 Port 1
	Eth1/11	40GbE	Cisco UCS C480-M5-7	VIC 1385 Port 1
	Eth1/12	40GbE	Cisco UCS C480-M5-8	VIC 1385 Port 1
	Eth1/13	40GbE	Cisco UCS C480-M5-9	VIC 1385 Port 1
	Eth1/14	40GbE	Cisco UCS C480-M5-10	VIC 1385 Port 1
	Eth1/15	40GbE	Cisco UCS C480-M5-11	VIC 1385 Port 1
	Eth1/16	40GbE	Cisco UCS C480-M5-12	VIC 1385 Port 1
	Eth1/17	40GbE	Cisco UCS C480-M5-13	VIC 1385 Port 1
	Eth1/18	40GbE	Cisco UCS C480-M5-14	VIC 1385 Port 1
	Eth1/19	40GbE	Cisco UCS C480-M5-15	VIC 1385 Port 1
	Eth1/20	40GbE	Cisco UCS C480-M5-16	VIC 1385 Port 1
	Eth1/21	40GbE	Cisco UCS C240-M5-1	VIC 1385 Port 1
	Eth1/22	40GbE	Cisco UCS C240-M5-2	VIC 1385 Port 1
	Eth1/23	40GbE	Cisco UCS C240-M5-3	VIC 1385 Port 1
	Eth1/24	40GbE	Cisco UCS C240-M5-4	VIC 1385 Port 1
	Eth1/25	40GbE	Cisco UCS C240-M5-5	VIC 1385 Port 1
	Eth1/26	40GbE	Cisco UCS C240-M5-6	VIC 1385 Port 1
	Eth1/27	40GbE	Cisco UCS C240-M5-7	VIC 1385 Port 1
	Eth1/28	40GbE	Cisco UCS C240-M5-8	VIC 1385 Port 1
	MGMT0	GbE	Customer's Management Switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1

Local Device	Local Port	Connection	Remote Device	Remote Port
	L2	GbE	Cisco UCS fabric interconnect B	L2

* The ports ETH1/3/1-4 are used with pair Nexus 9000 Switches design option.

Table 6 Cisco UCS Fabric Interconnect B - Cabling Information for Cisco UCS C480 M5 Server(s) Option

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3/1	40GbE	Cisco Nexus 3000 HA	Eth 1/7
	Eth1/3/2	QSFP to 4 SFP+ break-out cables	Cisco Nexus 3000 HA	Eth 1/8
	Eth1/3/3		Cisco Nexus 3000 HA	Eth 1/9
	Eth1/3/4		Cisco Nexus 3000 HA	Eth 1/10
	Eth1/3/1*		40GbE	Cisco Nexus 9000 A
	Eth1/3/2*	QSFP to 4 SFP+ break-out cables	Cisco Nexus 9000 A	Eth 1/6
	Eth1/3/3*		Cisco Nexus 9000 B	Eth 1/5
	Eth1/3/4*		Cisco Nexus 9000 B	Eth 1/6
	Eth1/4		40GbE	
	Eth1/5	40GbE	Cisco UCS C480-M5-1	VIC 1385 Port 2
	Eth1/6	40GbE	Cisco UCS C480-M5-2	VIC 1385 Port 2
	Eth1/7	40GbE	Cisco UCS C480-M5-3	VIC 1385 Port 2
	Eth1/8	40GbE	Cisco UCS C480-M5-4	VIC 1385 Port 2
	Eth1/9	40GbE	Cisco UCS C480-M5-5	VIC 1385 Port 2
	Eth1/10	40GbE	Cisco UCS C480-M5-6	VIC 1385 Port 2
	Eth1/11	40GbE	Cisco UCS C480-M5-7	VIC 1385 Port 2
	Eth1/12	40GbE	Cisco UCS C480-M5-8	VIC 1385 Port 2
	Eth1/13	40GbE	Cisco UCS C480-M5-9	VIC 1385 Port 2
Eth1/14	40GbE	Cisco UCS C480-M5-10	VIC 1385 Port 2	
Eth1/15	40GbE	Cisco UCS C480-M5-11	VIC 1385 Port 2	
Eth1/16	40GbE	Cisco UCS C480-M5-12	VIC 1385 Port 2	
Eth1/17	40GbE	Cisco UCS C480-M5-13	VIC 1385 Port 2	
Eth1/18	40GbE	Cisco UCS C480-M5-14	VIC 1385 Port 2	
Eth1/19	40GbE	Cisco UCS C480-M5-15	VIC 1385 Port 2	
Eth1/20	40GbE	Cisco UCS C480-M5-16	VIC 1385 Port 2	

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/21	40GbE	Cisco UCS C240-M5-1	VIC 1385 Port 2
	Eth1/22	40GbE	Cisco UCS C240-M5-2	VIC 1385 Port 2
	Eth1/23	40GbE	Cisco UCS C240-M5-3	VIC 1385 Port 2
	Eth1/24	40GbE	Cisco UCS C240-M5-4	VIC 1385 Port 2
	Eth1/25	40GbE	Cisco UCS C240-M5-5	VIC 1385 Port 2
	Eth1/26	40GbE	Cisco UCS C240-M5-6	VIC 1385 Port 2
	Eth1/27	40GbE	Cisco UCS C240-M5-7	VIC 1385 Port 2
	Eth1/28	40GbE	Cisco UCS C240-M5-8	VIC 1385 Port 2
	MGMT0	GbE	Customer's Management Switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1
	L2	GbE	Cisco UCS fabric interconnect A	L2

* The ports ETH1/3/1-4 are used with pair Nexus 9000 Switches design option.

Table 7 Cisco Nexus 3000 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 3000 HA	Eth1/3	10GbE	Cisco UCS fabric interconnect A	Eth1/3/1
	Eth1/4	10GbE	Cisco UCS fabric interconnect A	Eth1/3/2
	Eth1/5	10GbE	Cisco UCS fabric interconnect A	Eth1/3/3
	Eth1/6	10GbE	Cisco UCS fabric interconnect A	Eth1/3/4
	Eth1/7	10GbE	Cisco UCS fabric interconnect B	Eth1/3/1
	Eth1/8	10GbE	Cisco UCS fabric interconnect B	Eth1/3/2
	Eth1/9	10GbE	Cisco UCS fabric interconnect B	Eth1/3/3
	Eth1/10	10GbE	Cisco UCS fabric interconnect B	Eth1/3/4
	MGMT0	GbE	Customer's Management Switch	Any

Table 8 Cisco Nexus 9000-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9000 A	Eth1/1	10GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	10GbE	Uplink to Customer Data Switch B	Any
	Eth1/3	10GbE	Cisco UCS fabric interconnect A	Eth1/3/1
	Eth1/4	10GbE	Cisco UCS fabric interconnect A	Eth1/3/2
	Eth1/5	10GbE	Cisco UCS fabric interconnect B	Eth1/3/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/6	10GbE	Cisco UCS fabric interconnect B	Eth1/3/2
	Eth1/9*	10GbE	Cisco Nexus 9000 B	Eth1/9
	Eth1/10*	10GbE	Cisco Nexus 9000 B	Eth1/10
	Eth1/11*	10GbE	Cisco Nexus 9000 B	Eth1/11
	Eth1/12*	10GbE	Cisco Nexus 9000 B	Eth1/12
	MGMT0	GbE	Customer's Management Switch	Any

* The ports ETH1/9-12 can be replaced with E2/1 and E2/2 for 40G connectivity.



For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

Table 9 Cisco Nexus 9000-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9000 B	Eth1/1	10GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	10GbE	Uplink to Customer Data Switch B	Any
	Eth1/3	10GbE	Cisco UCS fabric interconnect A	Eth1/3/3
	Eth1/4	10GbE	Cisco UCS fabric interconnect A	Eth1/3/4
	Eth1/5	10GbE	Cisco UCS fabric interconnect B	Eth1/3/3
	Eth1/6	10GbE	Cisco UCS fabric interconnect B	Eth1/3/4
	Eth1/9*	10GbE	Cisco Nexus 9000 A	Eth1/9
	Eth1/10*	10GbE	Cisco Nexus 9000 A	Eth1/10
	Eth1/11*	10GbE	Cisco Nexus 9000 A	Eth1/11
	Eth1/12*	10GbE	Cisco Nexus 9000 A	Eth1/12
	MGMT0	GbE	Customer's Management Switch	Any

* The ports ETH1/9-12 can be replaced with E2/1 and E2/2 for 40G connectivity.



For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

Software Revisions

Table 10 details the software revisions used for validating various components of the Cisco UCS B/C480 M5 Scale-Out for SAP HANA.

Table 10 Hardware and Software Components of the Cisco UCS B/C480 M5 Scale-Out for SAP HANA

Vendor	Product	Version	Description
Cisco	Cisco UCSM	3.2(2d)	Cisco UCS Manager
Cisco	Cisco UCS 6332 UP FI	5.0(3)N2(3.22c)	Cisco UCS Fabric Interconnects
Cisco	Cisco UCS 5108 Blade Chassis	NA	Cisco UCS Blade Server Chassis
Cisco	Cisco UCS 2304 XP FEX	3.2(2c)	Cisco UCS Fabric Extenders for Blade Server chassis
Cisco	Cisco UCS B-Series M5 Servers	3.1(25d)	Cisco B-Series M5 Blade Servers
Cisco	Cisco UCS C-Series M5 Servers	3.1(2d) - CIMC Controller	Cisco C-Series M5 Blade Servers
Cisco	Cisco UCS C240 M5 Servers	3.1(2d) - CIMC Controller	Cisco C240 M5 Rack Servers for Management
Cisco	Cisco UCS VIC 1340/1380	4.2(2b)	Cisco UCS VIC 1240/1280 Adapters
Cisco	Cisco UCS VIC 1385	4.2(2b)	Cisco UCS VIC Adapter
Cisco	Cisco Nexus 9372PX Switches	7.0(3)I4(7)	Cisco Nexus 9372PX Switches
Cisco	Cisco Nexus 3524 Switches	7.0(3)I4(7)	Cisco Nexus N3K-C3524P Switch
RedHat	RedHat Enterprise Linux	7.4	Operating System to host SAP HANA
MapR	MapR Converged Enterprise Edition	6.0.0.20171109191718.GA	MapR Converged Enterprise Edition

Network Configuration

The following sections provide detailed procedures to configure the Cisco Nexus 3524 Switch for High Availability in the SAP HANA environment. The switch configuration in this section is based on the cabling plan described in the Device Cabling section. If the systems are connected on different ports, configure the switches accordingly by following the guidelines described below.

Cisco Nexus 3500 Series Switch Network Configuration

These steps provide the details for the initial Cisco Nexus 3524 Series Switch setup.

Cisco Nexus 3524 Initial Configuration

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



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

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

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

```
Do you want to enforce secure password standard (yes/no): yes
```

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

```
---- Basic System Configuration Dialog ----
```

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

```
Please register Cisco Nexus 3500 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus devices must be registered to receive entitled
support services.
```

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

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

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

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

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

```
Enter the switch name : <<var_nexus_HA_hostname>>
```

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

```
Mgmt0 IPv4 address : <<var_nexus_HA_mgmt0_ip>>
```

```

Mgmt0 IPv4 netmask : <<var_nexus_HA_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]:

  IPv4 address of the default gateway : <<var_nexus_HA_mgmt0_gw>>
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:

  Type of ssh key you would like to generate (dsa/rsa) : rsa
  Number of key bits <768-2048> : 2048
Configure the ntp server? (yes/no) [n]: y

  NTP server IPv4 address : <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut) [noshut]:
Configure CoPP System Policy Profile ( default / 12 / 13 ) [default]:

The following configuration will be applied:
  switchname <<var_nexus_HA_hostname>>
  interface mgmt0
  ip address <<var_nexus_HA_mgmt0_ip>> <<var_nexus_HA_mgmt0_netmask>>
  vrf context management
  ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
  no shutdown
  no telnet server enable
  ssh key rsa 2048 force
  ssh server enable
  ntp server <<var_global_ntp_server_ip>>
  system default switchport
  no system default switchport shutdown
  policy-map type control-plane copp-system-policy ( default )

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

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

[#####] 100%
Copy complete, now saving to disk (please wait)...

```

Enable Appropriate Cisco Nexus 3524 Switch Features and Settings

The following commands enable IP switching feature and set default spanning tree behaviors:

1. On each Nexus 3524, enter configuration mode:

```
config terminal
```

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

```
feature lacp
feature interface-vlan
feature lldp
```

3. Save the running configuration to start-up:

```
copy run start
```

Create Global Policy to Enable Jumbo Frame and Apply the Policy to System Wide

```

policy-map type network-qos jumbo
  class type network-qos class-default

    mtu 9216

system qos
  service-policy type network-qos jumbo

```

Create VLANs for SAP HANA Traffic

To create the necessary VLANs, complete the following step on both switches:

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

```

vlan <<var_storage_vlan_id>>
name HANA-Storage

vlan <<var_internal_vlan_id>>
name HANA-Internal

vlan <<var_inband_vlan_id>>
name NFS-IPMI

vlan <<var_mapr-01_vlan_id>>
name MapR-01

vlan <<var_mapr-02_vlan_id>>
name MapR-02

vlan <<var_mapr-03_vlan_id>>
name MapR-03

```

Configure Network Interfaces Connecting to Cisco UCS Fabric Interconnect

1. Define a port description for the interface connecting to <<var_ucs_clustername>>-A.

```

interface Eth1/3-6
description <<var_ucs_clustername>>-A:1/3

```

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

```

interface eth1/3-6
channel-group 3 mode active
no shutdown

```

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

```

interface Po3
description <<var_ucs_clustername>>-A

```

4. Make the port channel a switchport, and configure a trunk to allow internal HANA VLANs.

```

switchport
switchport mode trunk
switchport trunk allowed vlan <<var_storage_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_id>>,<<var_mapr-01_vlan_id>>,<<var_mapr-02_vlan_id>>,<<var_mapr-03_vlan_id>>

```

5. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree bpduguard enable
```

6. Bring up the port-channel.

```
no shutdown
```

7. Define a port description for the interface connecting to <<var_ucs_clustername>>-B.

```
interface Eth1/7-10
description <<var_ucs_clustername>>-B:1/3
```

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

```
interface Eth1/7-10
channel-group 4 mode active
no shutdown
```

9. Define a description for the port-channel connecting to <<var_ucs_clustername>>-B.

```
interface Po4
description <<var_ucs_clustername>>-B
```

10. Make the port-channel a switchport, and configure a trunk to allow internal HANA VLANs.

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_storage_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_id>>,
<<var_mapr-01_vlan_id>>,<<var_mapr-02_vlan_id>>,<<var_mapr-03_vlan_id>>
```

11. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree bpduguard enable
```

12. Bring up the port channel.

```
no shutdown
```

13. Save the running configuration to start-up configuration.

```
copy run start
```

Cisco Nexus 9000 Series Switch Network Configuration

This section provides the detailed procedure to configure the Cisco Nexus 9000 Switches for SAP HANA environment. The switch configuration in this section is based on the cabling plan described in section [Physical Device Cabling](#); if the systems are connected on different ports, configure the switches according to the guidelines described in this section.



The configuration steps detailed in this section provides guidance for configuring the Cisco Nexus 9000 running release 7.0(3)I5(1) within a multi-VDC environment.

Cisco Nexus 9000 A Initial Configuration

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



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

```

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

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

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

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

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

Do you want to enforce secure password standard (yes/no) [y]:
  Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
  Enter the switch name : <<var_nexus_A_hostname>>
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
    Mgmt0 IPv4 address : <<var_nexus_A_mgmt0_ip>>
    Mgmt0 IPv4 netmask : <<var_nexus_A_mgmt0_netmask>>
  Configure the default gateway? (yes/no) [y]:
    IPv4 address of the default gateway : <<var_nexus_A_mgmt0_gw>>
  Configure advanced IP options? (yes/no) [n]:
  Enable the telnet service? (yes/no) [n]:
  Enable the ssh service? (yes/no) [y]:
    Type of ssh key you would like to generate (dsa/rsa) [rsa]:
    Number of rsa key bits <1024-2048> [2048]:
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : <<var_global_ntp_server_ip>>
  Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:

The following configuration will be applied:
  password strength-check
  switchname <<var_nexus_A_hostname>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 2048 force
feature ssh

```

```

ntp server <<var_global_ntp_server_ip>>
copp profile strict
interface mgmt0
ip address <<var_nexus_A_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown

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

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

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

```

Cisco Nexus 9000 B Initial Configuration

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



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

```

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

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

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

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

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

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

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

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

Enter the switch name : <<var_nexus_B_hostname>>

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

Mgmt0 IPv4 address : <<var_nexus_B_mgmt0_ip>>

Mgmt0 IPv4 netmask : <<var_nexus_B_mgmt0_netmask>>

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

IPv4 address of the default gateway : <<var_nexus_B_mgmt0_gw>>

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

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

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

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

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

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

NTP server IPv4 address : <<var_global_ntp_server_ip>>

```

```

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

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

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

The following configuration will be applied:
  password strength-check
  switchname <<var_nexus_B_hostname>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
  no feature telnet
  ssh key rsa 2048 force
  feature ssh
  ntp server <<var_global_ntp_server_ip>>
  copp profile strict
interface mgmt0
ip address <<var_nexus_B_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown

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

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

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

```

Enable Appropriate Cisco Nexus 9000 Series Switches—Features and Settings

Cisco Nexus 9000 A and Cisco Nexus 9000 B

The following commands enable the IP switching feature and set the default spanning tree behaviors:

1. On each Nexus 9000, enter configuration mode:

```
config terminal
```

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

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

3. Configure spanning tree defaults:

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

4. Save the running configuration to start-up:

```
copy run start
```

Create VLANs for SAP HANA Traffic

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To create the necessary VLANs, complete the following step on both switches:

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

```
vlan <<var_storage_vlan_id>>
name HANA-Storage

vlan <<var_mgmt_vlan_id>>
name Management

vlan <<var_internal_vlan_id>>
name HANA-Internal

vlan <<var_backup_vlan_id>>
name HANA-Backup

vlan <<var_client_vlan_id>>
name HANA-Client

vlan <<var_appserver_vlan_id>>
name HANA-AppServer

vlan <<var_datasource_vlan_id>>
name HANA-DataSource

vlan <<var_replication_vlan_id>>
name HANA-Replication

vlan <<var_inband_vlan_id>>
name NFS-IPMI

vlan <<var_mapr-01_vlan_id>>
name MapR-01

vlan <<var_mapr-02_vlan_id>>
name MapR-02

vlan <<var_mapr-03_vlan_id>>
name MapR-03
```

Configure Virtual Port-Channel Domain

Cisco Nexus 9000 A

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

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

```
vpc domain <<var_nexus_vpc_domain_id>>
```

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

```
role priority 10
```

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

```
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source <<var_nexus_A_mgmt0_ip>>
```

4. Enable following features for this vPC domain:

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

Cisco Nexus 9000 B

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

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

```
vpc domain <<var_nexus_vpc_domain_id>>
```

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

```
role priority 20
```

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

```
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source <<var_nexus_B_mgmt0_ip>>
```

4. Enable following features for this vPC domain:

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

Configure Network Interfaces for the VPC Peer Links

Cisco Nexus 9000 A

1. Define a port description for the interfaces connecting to VPC Peer <<var_nexus_B_hostname>>.

```
interface Eth2/1
description VPC Peer <<var_nexus_B_hostname>>:2/1

interface Eth2/2
description VPC Peer <<var_nexus_B_hostname>>:2/2
```

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

```
interface Eth2/1-2
channel-group 1 mode active
no shutdown
```

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

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

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

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_storage_vlan_id>>,<<var_mgmt_vlan_id>>,<<var_boot_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_
id>>,<<var_backup_vlan_id>>,<<var_client_vlan_id>>,<<var_appserver_vlan_id>>,<<var_datasource_vlan_id>>,<<var_replication_vlan_id>>
```

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

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

Cisco Nexus 9000 B

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

```
interface Eth2/1
description VPC Peer <<var_nexus_A_hostname>>:2/1

interface Eth2/2
description VPC Peer <<var_nexus_A_hostname>>:2/2
```

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

```
interface Eth2/1-2
channel-group 1 mode active
no shutdown
```

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

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

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

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_storage_vlan_id>>,<<var_mgmt_vlan_id>>,<<var_boot_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_
id>>,<<var_backup_vlan_id>>,<<var_client_vlan_id>>,<<var_appserver_vlan_id>>,<<var_datasource_vlan_id>>,<<var_replication_vlan_id>>
```

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

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

Configure Network Interfaces with Cisco UCS Fabric Interconnect

Cisco Nexus 9000 A and Cisco Nexus 9000 B

1. Define a port description for the interface connecting to <<var_ucs_clustername>>-A.

```
interface Eth1/3-4
description <<var_ucs_clustername>>-A:1/3
```

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

```
interface eth1/3-4
channel-group 11 mode active
no shutdown
```

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

```
interface Po11
description <<var_ucs_clustername>>-A
```

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

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_storage_vlan_id>>,<<var_mgmt_vlan_id>>,<<var_boot_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_
id>>,<<var_backup_vlan_id>>, <<var_client_vlan_id>>, <<var_appserver_vlan_id>>, <<var_datasource_vlan_id>>,
<<var_replication_vlan_id>>
```

5. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

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

```
mtu 9216
```

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

```
vpc 11
no shutdown
```

8. Define a port description for the interface connecting to <<var_ucs_clustername>>-B.

```
interface Eth1/5-6
description <<var_ucs_clustername>>-B:1/3
```

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

```
interface Eth1/5-6
channel-group 12 mode active
no shutdown
```

10. Define a description for the port-channel connecting to <<var_ucs_clustername>>-B.

```
interface Po12
description <<var_ucs_clustername>>-B
```

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

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_storage_vlan_id>>,<<var_mgmt_vlan_id>>,<<var_boot_vlan_id>>,<<var_internal_vlan_id>>,<<var_inband_vlan_
id>>,<<var_backup_vlan_id>>, <<var_client_vlan_id>>, <<var_appserver_vlan_id>>, <<var_datasource_vlan_id>>,
<<var_replication_vlan_id>>
```

12. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

13. 53. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

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

```
vpc 12
no shutdown
```

(Optional) Configure Network Interfaces for SAP HANA Backup/Data Source/Replication

You can define the port channel for each type of network to have a dedicated bandwidth. Below is an example to create a port channel for Backup Network; these cables are connected to Storage for Backup. The following example assumes two ports (Ethernet 1/29 and 1/30) are connected to a dedicated NFS Storage to backup SAP HANA.

Cisco Nexus 9000 A and Cisco Nexus 9000 B

1. Define a port description for the interface connecting to <<var_node01>>.

```
interface Eth1/29
description <<var_backup_node01>>:<<Port_Number>>
```

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

```
interface eth1/29
channel-group 21 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <<var_backup_node01>>.

```
interface Po21
description <<var_backup_vlan_id>>
```

4. Make the port-channel a switchport, and configure a trunk to allow NFS VLAN for DATA.

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_backup_vlan_id>>
```

5. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

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

```
mtu 9216
```

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

```
vpc 21
no shutdown
```

- Define a port description for the interface connecting to <<var_node02>>.

```
interface Eth1/30
description <<var_backup_node01>>:<<Port_Number>>
```

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

```
channel-group 22 mode active
no shutdown
```

- Define a description for the port-channel connecting to <<var_node02>>.

```
interface Po22
description <<var_backup_node02>>
```

- Make the port channel a switchport, and configure a trunk to allow NFS VLAN for DATA.

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_backup_vlan_id>>
```

- Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

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

```
mtu 9216
```

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

```
vpc 22
no shutdown
```

(Optional) Management Plane Access for Cisco UCS Servers

This is an optional step, which can be used to implement a management plane access for the Cisco UCS servers.

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To enable management access across the IP switching environment, complete the following steps:



You may want to create a dedicated Switch Virtual Interface (SVI) on the Nexus data plane to test and troubleshoot the management plane. If an L3 interface is deployed be sure it is deployed on both Cisco Nexus 9000s to ensure Type-2 VPC consistency.

- Define a port description for the interface connecting to the management plane.

```
interface Eth1/<<interface_for_in_band_mgmt>>
description IB-Mgmt:<<mgmt_uplink_port>>
```

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

```
channel-group 6 mode active
no shutdown
```

3. Define a description for the port-channel connecting to management switch.

```
interface Po6
description IB-Mgmt
```

4. Configure the port as an access VLAN carrying the InBand management VLAN traffic.

```
switchport
switchport mode access
switchport access vlan <<var_inband_vlan_id>>
```

5. Make the port channel and associated interfaces normal spanning tree ports.

```
spanning-tree port type normal
```

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

```
vpc 6
no shutdown
```

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

```
copy run start
```

Uplink into Existing Network Infrastructure

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

Cisco UCS Configuration

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

Initial Setup of Cisco UCS 6332 Fabric Interconnect

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

Cisco UCS 6332 Fabric Interconnect A

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

1. Connect to the console port on the first Cisco UCS 6300 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": <<var_password>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)?
(yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter

```

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

Cisco UCS 6332 Fabric Interconnect B

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

1. Connect to the console port on the second Cisco UCS 6332 Fabric Interconnect.

```

Enter the configuration method: console
Installer has detected the presence of a peer Fabric interconnect. This Fabric
interconnect will be added to the cluster. Do you want to continue {y|n}? y
Enter the admin password for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no):
y

```


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

Cisco UCS for SAP HANA

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6332 Fabric Interconnect cluster address.
2. Click the Launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to the Cisco UCS Manager.



Upgrade Cisco UCS Manager Software to Version 3.2(2d)

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

Add Block of IP Addresses for KVM Access

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



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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Pools > root > IP Pools > IP Pool ext-mgmt.
3. In the Actions pane, choose Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.
5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click the Admin tab in the navigation pane.
2. Choose All > Timezone Management.
3. In the Properties pane, choose the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter <<var_global_ntp_server_ip>> and click OK.
7. Click OK.

Cisco UCS Blade Chassis Connection Options

For the Cisco UCS 6300 Series Fabric Extenders, two configuration options are available: pinning and portchannel. 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.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies adding the Cisco UCS B-Series chassis and additional fabric extenders for C-Series connectivity.

To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and choose 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 server ports that are connected between the fabric extenders (FEXes) within the chassis and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.
5. Click Save Changes.
6. Click OK.

Equipment

Main Topology View

Fabric Interconnects

Servers

Thermal

Global Policies

Autoconfig Policies

Server Inheritance Policies

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, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Choose Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Choose the ports that are connected to the chassis and / or to the Cisco C-Series Server (two per FI), right-click them, and choose Configure as Server Port.
5. Click Yes to confirm server ports and click OK.

- Verify that the ports connected to the chassis and / or to the Cisco C-Series Server are now configured as server ports.

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

Ethernet Ports

Advanced Filter Export Print All Unconfigured Network Server FCoE Uplink Unified Uplink Appliance Storage FCoE Storage » ⚙

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	17	8C:60:4F:BC:C5:34	Server	Physical	↑ Up	↑ Enabled
1	0	18	8C:60:4F:BC:C5:38	Server	Physical	↑ Up	↑ Enabled
1	0	19	8C:60:4F:BC:C5:3C	Server	Physical	↑ Up	↑ Enabled
1	0	20	8C:60:4F:BC:C5:40	Server	Physical	↑ Up	↑ Enabled
1	0	21	8C:60:4F:BC:C5:44	Server	Physical	↑ Up	↑ Enabled
1	0	22	8C:60:4F:BC:C5:48	Server	Physical	↑ Up	↑ Enabled
1	0	23	8C:60:4F:BC:C5:4C	Server	Physical	↑ Up	↑ Enabled
1	0	24	8C:60:4F:BC:C5:50	Server	Physical	↑ Up	↑ Enabled
1	0	26	8C:60:4F:BC:C5:58	Server	Physical	↑ Up	↑ Enabled
1	0	27	8C:60:4F:BC:C5:5C	Server	Physical	↑ Up	↑ Enabled
1	0	28	8C:60:4F:BC:C5:60	Server	Physical	↑ Up	↑ Enabled
1	0	29	8C:60:4F:BC:C5:64	Server	Physical	↑ Up	↑ Enabled
1	0	32	8C:60:4F:BC:C5:70	Server	Physical	↑ Up	↑ Enabled
1	0	33	8C:60:4F:BC:C5:74	Server	Physical	↑ Up	↑ Enabled

- Choose ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
- Click Yes to confirm uplink ports and click OK.
- Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
- Expand Ethernet Ports.
- Choose the ports that are connected to the chassis or to the Cisco C-Series Server (two per FI), right-click them, and choose Configure as Server Port.
- Click Yes to confirm server ports and click OK.
- Choose ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
- Click Yes to confirm the uplink ports and click OK.

Configure Breakout Ports

The 40 GB Ethernet ports on Cisco UCS 6300 Fabric Interconnects can be configured as four breakout 10 GB ports using a supported breakout cable.



Configuring breakout ports requires rebooting the Fabric Interconnect. Any existing configuration on a port is erased. It is recommended to break out all required ports in a single transaction.

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Choose Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Choose ports that are connected to the Cisco Nexus 3524 switches, right-click them, and choose Configure Breakout Port.
5. Click Yes to confirm the Breakout ports and click OK.
6. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
7. Expand Ethernet Ports.
8. Choose ports that are connected to the Cisco Nexus 3524 switches, right-click them, and choose Configure Breakout Port.
9. Click Yes to confirm the Breakout ports and click OK.



When you configure a breakout port, you can configure each 10 GB sub-port as server, uplink, FCoE uplink, FCoE storage or appliance as required. Unified Ports cannot be configured as Breakout Ports.

Acknowledge Cisco UCS Chassis and Rack-Mount Servers

To acknowledge all Cisco UCS chassis and Rack Mount Servers, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Expand Chassis and choose each chassis that is listed.
3. Right-click each chassis and choose 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 choose Acknowledge Server.
7. Click Yes and then click OK to complete acknowledging the Rack Mount Servers.

Create Uplink Port Channels

A separate uplink port-channel for each of the network zones are defined as per SAP. For example, create port-channel 11 on fabric interconnect A and port-channel 12 on fabric interconnect B for Client zone network.

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
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. Choose Create Port Channel.
6. Enter 11 as the unique ID of the port channel.
7. Enter vPC-11-Nexus as the name of the port channel.
8. Click Next.



9. Choose the following ports to be added to the port channel:
 - Slot ID 1 and port 1
 - Slot ID 1 and port 2
10. If the breakout cables are used for Uplink connectivity. Choose the following ports to be added to the port channel:
 - Slot ID 1, Aggregated Port ID 1 and port 1
 - Slot ID 1, Aggregated Port ID 1 and port 2
 - Slot ID 1, Aggregated Port ID 1 and port 3
 - Slot ID 1, Aggregated Port ID 1 and port 4
 - Slot ID 1, Aggregated Port ID 2 and port 1

- Slot ID 1, Aggregated Port ID 2 and port 2
 - Slot ID 1, Aggregated Port ID 2 and port 3
 - Slot ID 1, Aggregated Port ID 2 and port 4
11. Click >> to add the ports to the port channel.
 12. Click Finish to create the port channel.
 13. Click OK.
 14. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
 15. Right-click Port Channels.
 16. Choose Create Port Channel.
 17. Enter 12 as the unique ID of the port channel.
 18. Enter vPC-12-NEXUS as the name of the port channel.
 19. Click Next.
 20. Choose the following ports to be added to the port channel:
 - Slot ID 1 and port 1
 - Slot ID 1 and port 2
 21. If the breakout cables are used for Uplink connectivity. Choose the following ports to be added to the port channel:
 - Slot ID 1, Aggregated Port ID 1 and port 1
 - Slot ID 1, Aggregated Port ID 1 and port 2
 - Slot ID 1, Aggregated Port ID 1 and port 3
 - Slot ID 1, Aggregated Port ID 1 and port 4
 - Slot ID 1, Aggregated Port ID 2 and port 1
 - Slot ID 1, Aggregated Port ID 2 and port 2
 - Slot ID 1, Aggregated Port ID 2 and port 3
 - Slot ID 1, Aggregated Port ID 2 and port 4
 22. Click >> to add the ports to the port channel.
 23. Click Finish to create the port channel.

24. Click OK.

Repeat steps 1-24 to create Additional port-channel for each network zone based on your Data Center requirements.

If you are using single Nexus 3524 for IOM failover, create port-channel 5 on fabric interconnect A and portchannel 6 on fabric interconnect B for Internal network zone.

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
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. Choose Create Port Channel.
6. Enter 5 as the unique ID of the port channel.
7. Enter N3k-Uplink as the name of the port channel.
8. Click Next.
9. Choose the following ports to be added to the port channel:
 - Slot ID 1, Aggregated Port ID 3 and port 1
 - Slot ID 1, Aggregated Port ID 3 and port 2
 - Slot ID 1, Aggregated Port ID 3 and port 3
 - Slot ID 1, Aggregated Port ID 3 and port 4
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.
14. Right-click Port Channels.
15. Choose Create Port Channel.
16. Enter 6 as the unique ID of the port channel.
17. Enter N3k-Uplink as the name of the port channel.

18. Click Next.
19. Choose the following ports to be added to the port channel:
 - Slot ID 1, Aggregated Port ID 3 and port 1
 - Slot ID 1, Aggregated Port ID 3 and port 2
 - Slot ID 1, Aggregated Port ID 3 and port 3
 - Slot ID 1, Aggregated Port ID 3 and port 4
20. Click >> to add the ports to the port channel.
21. Click Finish to create the port channel.
22. Click OK.

Create New Organization

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

To create organization unit, complete the following steps:

1. In Cisco UCS Manager, on the Tool bar click New.
2. From the drop-down list choose Create Organization.
3. Enter the Name as HANA.
4. (Optional) Enter the Description as Org for HANA.
5. Click OK to create the Organization.

Create MAC Address Pools

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Pools > root > Sub-Organization > HANA.
3. In this procedure, two MAC address pools are created, one for each switching fabric.
4. Right-click MAC Pools under the root organization.
5. Choose Create MAC Pool to create the MAC address pool.
6. Enter FI-A as the name of the MAC pool.
7. (Optional) Enter a description for the MAC pool.

8. Choose Assignment Order Sequential.
9. Click Next.
10. Click Add.
11. Specify a starting MAC address.
12. The recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as Fabric Interconnect A addresses.
13. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

Create a Block of MAC Addresses



First MAC Address : Size :

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

00:25:B5:xx:xx:xx



14. Click OK.
15. Click Finish.
16. In the confirmation message, click OK.
17. Right-click MAC Pools under the HANA organization.
18. Choose Create MAC Pool to create the MAC address pool.
19. Enter FI-B as the name of the MAC pool.
20. (Optional) Enter a description for the MAC pool.
21. Click Next.
22. Click Add.

23. Specify a starting MAC address.



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

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

25. Click OK.

26. Click Finish.

27. In the confirmation message, click OK.



You can also define separate MAC address Pool for each Network Zone. Follow the above steps 1-16 to create MAC address pool for each Network Zone.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Pools > root.
3. Right-click UUID Suffix Pools.
4. Choose 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. Choose Sequential for Assignment Order
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.

15. Click OK.

Power Policy

To run Cisco UCS with two independent power distribution units, the redundancy must be configured as Grid. Complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and choose Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Redundancy field in Power Policy to Grid.
4. Click Save Changes.
5. Click OK.

Power Control Policy

The Power Capping feature in Cisco UCS is designed to save power with legacy data center use cases. This **feature does not contribute 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 make sure there is sufficient power supply for high performance and critical applications such as SAP HANA.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Right-click Power Control Policies.
4. Choose 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



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.

OK

Cancel

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Right-click Host Firmware Packages.

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

Create Host Firmware Package

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:

<input type="checkbox"/>	Adapter
<input type="checkbox"/>	BIOS
<input type="checkbox"/>	Board Controller
<input type="checkbox"/>	CIMC
<input type="checkbox"/>	FC Adapters
<input type="checkbox"/>	Flex Flash Controller
<input type="checkbox"/>	GPUs
<input type="checkbox"/>	HBA Option ROM
<input type="checkbox"/>	Host NIC
<input type="checkbox"/>	Host NIC Option ROM

OK

Cancel

Create Local Disk Configuration Policy

A local disk configuration policy configures SAS local drives that have been installed on a server through the onboard RAID controller of the local drive.

To create a local disk configuration policy for HANA servers for Local OS disks, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Right-click Local Disk Config Policies.

4. Choose Create Local Disk Configuration Policy.
5. Enter RAID1 as the local disk configuration policy name.
6. Change the mode to RAID 1.
7. Click OK to create the local disk configuration policy.

Create Local Disk Configuration Policy



Name :

Description :

Mode :

Protect Configuration :

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

FlexFlash

FlexFlash State : Disable 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

OK

Cancel

8. Click OK.

To create a local disk configuration policy for MapR servers, complete the following steps:

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

2. Choose Policies > root.
3. Right-click Local Disk Config Policies.
4. Choose Create Local Disk Configuration Policy.
5. Enter MapR as the local disk configuration policy name.
6. Change the mode to Any Configuration.
7. Click OK to create the local disk configuration policy.

Create Server BIOS Policy

To get the best performance for HANA it is required to configure the Server BIOS accurately. To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click BIOS Policies.
4. Choose Create BIOS Policy.
5. Enter HANA-BIOS as the BIOS policy name.
6. Click OK.

Create BIOS Policy



Name :

Description :

Reboot on BIOS Settings Change :

OK

Cancel

7. Under BIOS Policies, click the newly created HANA-BIOS Policy.
8. In the Main pane, under BIOS Setting choose Disabled for Quiet Boot.
9. Click the Advance tab.
10. The recommendation from SAP for SAP HANA is to disable all Processor C States. This will force the CPU to stay on maximum frequency and allow SAP HANA to run with best performance.
11. Under Processor choose Disabled for all C-States.
12. Set HPC for CPU Performance, Performance for Power Technology, Energy Performance.

Servers / Policies / root / BIOS Policies / HANA-BIOS

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Performance
Energy Performance	Performance
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default

+ Add - Delete Info

Save Changes Reset Values

13. Click RAS Memory.

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

Servers / Policies / root / BIOS Policies / HANA-BIOS

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Enabled
Memory RAS configuration	Maximum Performance

15. Click Serial Port.

16. Choose Enabled for Serial Port A enable.

Servers / Policies / root / BIOS Policies / HANA-BIOS

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
Serial port A enable	Enabled

17. Click Server Management.

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

Servers / Policies / root / BIOS Policies / HANA-BIOS

Main Advanced Boot Options **Server Management** Events

Advanced Filter Export Print

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

19. Click Save Changes.

To create a server BIOS policy for MapR server, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click BIOS Policies.
4. Choose Create BIOS Policy.
5. Enter MapR-BIOS as the BIOS policy name.

6. Click OK.
7. Under BIOS Policies, Click on newly created MapR-BIOS Policy.
8. In the Main pane, under BIOS Setting choose Disabled for Quiet Boot.
9. Click Advance tab.
10. Under Processor choose Disabled for all C-States. This will force the CPU to stay on maximum frequency and allow MapR Server to run with best performance.
11. Choose Enabled for Enhanced Intel SpeedStep Tech, Intel Hyper Threading Tech, Intel Turbo Boost Tech.
12. Set High Throughput for CPU Performance.
13. Choose Performance for Power Technology, Energy Performance, DRAM Clock Throttling.

Servers / Policies / root / BIOS Policies / MapR-BIOS

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Platform Default
CPU Performance	High Throughput
Core Multi Processing	All
DRAM Clock Throttling	Performance
Direct Cache Access	Platform Default
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Enabled
Execute Disable Bit	Platform Default
Frequency Floor Override	Platform Default
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	Enabled
Intel Virtualization Technology	Disabled

Add Delete Info

14. Click RAS Memory.
15. Choose Maximum-Performance for Memory RAS Configuration and Enabled for NUMA optimized.
16. 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 in case the server hangs or if the Linux kernel crashes, when the dump is required. Configure the speed in the Server Management tab of the BIOS Policy. To create the serial over LAN policy, complete the following steps:

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

Create Serial over LAN Policy



Name	:	<input type="text" value="SoL-Console"/>
Description	:	<input type="text"/>
Serial over LAN State	:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Speed	:	<input type="text" value="115200"/>

OK **Cancel**

Update Default Maintenance Policy

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Choose 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	
Delete		Name	: default
Show Policy Usage		Description	: <input type="text"/>
Use Global		Owner	: Local
		Soft Shutdown Timer	: 150 Secs <input type="text"/>
		Storage Config. Deployment Policy	: <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack
		Reboot Policy	: <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic
		<input checked="" type="checkbox"/> On Next Boot	(Apply pending changes at next reboot.)

IPMI Access Profiles

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

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

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

4. Choose Create IPMI Access Profile.
5. Enter HANA-IPMI as the Profile name.
6. Click + button.
7. Enter Username in the Name field and password.
8. Choose Admin as Role.

Create IPMI User



Name :

Password :

Confirm Password :

Role : Read Only Admin

Description :



9. Click OK to create user.
10. Click OK to Create IPMI Access Profile.

Create IPMI Access Profile



Name :

Description :

IPMI Over LAN : Disable Enable

IPMI Users

+ - Advanced Filter ↑ Export ↑ Print ⚙	
Name	Role
sapadm	Admin

+ Add 🗑 Delete ℹ Info

OK
Cancel

Network Configuration

The core network requirements for SAP HANA are managed by the Cisco UCS defaults. Cisco UCS is based on 40GbE and provides redundancy through 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 inter-node flows on FI A and the Storage traffic is on FI B. The inter-node traffic flows from a Blade Server to the Fabric Interconnect A and back to other Blade Server. The storage traffic flows from a HANA Server to the Fabric Interconnect B and back to MapR Server.

Set Jumbo Frames in Cisco UCS Fabric

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

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

2. Choose LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the MTU Column, enter 9216 in the box.
5. Check Enabled Under Priority for Silver.
6. Click Save Changes in the bottom of the window.
7. Click OK.

LAN

LAN Uplinks VLANs Server Links MAC Identity Assignment IP Identity Assignment **QoS** Global Policies Faults Events FSM

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	9128	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	9128	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	5	N/A	9128	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	9128	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	100	9128	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	none	N/A	fc	N/A

Network Control Policy

Update Default Network Control Policy to Enable CDP

CDP needs to be enabled to learn the MAC address of the End Point. To update default Network Control Policy, complete the following steps:

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

LAN / Policies / root / Network Control Policies / default

General	Events
Actions <hr/> Delete Show Policy Usage Use Global	Properties <hr/> Name : default Description : <input type="text"/> Owner : Local CDP : <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled MAC Register Mode : <input checked="" type="radio"/> Only Native Vlan <input type="radio"/> All Host Vlans Action on Uplink Fail : <input checked="" type="radio"/> Link Down <input type="radio"/> Warning MAC Security <hr/> Forge : <input checked="" type="radio"/> Allow <input type="radio"/> Deny LLDP <hr/> Transmit : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled Receive : <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

Create Network Control Policy for Internal Network

In order to keep the vNIC links up in case of a Nexus 3524 failure, create the Network Control Policy for Internal Network. To create Network Control Policy for Internal Network, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose LAN > Policies > root > Network Control Policies > right-click and choose Create Network Control Policy.
3. Enter Internal as the Name of the Policy.
4. For CDP, choose Enabled radio button.
5. For Action on Uplink Fail, choose Warning radio button.
6. Click OK.

Create Network Control Policy



Name :

Description :

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

Warning

IMPORTANT: If the Action on Uplink Fail is set to **Warning**, the fabric **will not** fail over if uplink connectivity is lost

MAC Security

All Disabled

OK

Cancel

LAN Configurations

Within Cisco UCS, all the network types for a SAP HANA system are defined VLANs. The 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 221 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 need not be created.

Create VLANs

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

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



In this procedure, nine VLANs are created.

2. Choose LAN > LAN Cloud.
3. Right-click VLANs.

4. Choose Create VLANs.
5. Enter HANA-Internal as the name of the VLAN to be used for HANA Node to Node network.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter <<var_internal_vlan_id>> as the ID of the HANA Node to Node network.
8. Keep the Sharing Type as None.
9. Click OK.

Create VLANs

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

10. Repeat the Steps 1-9 for each VLAN.
11. Create VLAN for HANA-AppServer.

Create VLANs

VLAN Name/Prefix : HANA-AppServer

Multicast Policy Name : <not set>

[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 : 223

Sharing Type : None Primary Isolated Community

12. Create VLAN for HANA-Backup.

Create VLANs

VLAN Name/Prefix : HANA-Backup

Multicast Policy Name : <not set>

[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 : 221

Sharing Type : None Primary Isolated Community

13. Create VLAN for HANA-Client.

Create VLANs

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

14. Create VLAN for HANA-DataSource.

Create VLANs

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

15. Create VLAN for HANA-Replication.

Create VLANs

VLAN Name/Prefix : HANA-Replication

Multicast Policy Name : <not set>

[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 : 225

Sharing Type : None Primary Isolated Community

16. Create VLAN for HANA-Storage.

Create VLANs

VLAN Name/Prefix : HANA-Storage

Multicast Policy Name : <not set>

[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 : 110

Sharing Type : None Primary Isolated Community

17. Create VLAN for NFS-IPMI for Inband Access.

Create VLANs

VLAN Name/Prefix : NFS-IPMI

Multicast Policy Name : <not set>

[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 : 197

Sharing Type : None Primary Isolated Community

18. Create VLAN for Management.

Create VLANs

VLAN Name/Prefix : Management

Multicast Policy Name : <not set>

[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 : 196

Sharing Type : None Primary Isolated Community

19. Create VLAN for MapR-01.

Create VLANs

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

20. Create VLAN for MapR-02.

Create VLANs

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

Create VLAN Groups

For easier management and bandwidth allocation to a dedicated uplink on the Fabric Interconnect, VLAN Groups are created within the Cisco UCS. SAP HANA uses the following VLAN groups:

- Admin Zone
- Client Zone
- Internal Zone
- Backup Network
- Replication Network

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

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



In this procedure, five VLAN Groups are created. Based on the solution requirement create VLAN groups, it not required to create all five VLAN groups.

2. Choose LAN > LAN Cloud.
3. Right-click VLAN Groups.
4. Choose Create VLAN Groups.
5. Enter Admin-Zone as the name of the VLAN Group used for Infrastructure network.
6. Choose Management.

1 Select VLANs

2 Add Uplink Ports

3 Add Port Channels

Create VLAN Group

Name :

VLANs

Advanced Filter Export Print No Native VLAN

Select	Name	Native VLAN
<input type="checkbox"/>	HANA-Client	<input type="radio"/>
<input type="checkbox"/>	HANA-DataSource	<input type="radio"/>
<input type="checkbox"/>	HANA-Internal	<input type="radio"/>
<input type="checkbox"/>	HANA-Replication	<input type="radio"/>
<input type="checkbox"/>	HANA-Storage	<input type="radio"/>
<input checked="" type="checkbox"/>	Management	<input type="radio"/>

Create VLAN

< Prev Next > Finish Cancel

7. Click Next.

8. Click Next on Add Uplink Ports.

9. Choose the port-channels created for the Admin Network.

1 Select VLANs

2 Add Uplink Ports

3 Add Port Channels

Create VLAN Group

Port Channels		
Name	Fabric ID	ID
N3k-U...	A	5
N3k-U...	B	6

>>
<<

Selected Port Channels		
Name	Fabric ID	ID
vPC-31-Nexus	A	31
vPC-32-Nexus	B	32

< Prev Next > **Finish** Cancel

10. Click Finish.

11. Follow the steps 1-10 for each VLAN Group.

12. Choose Create VLAN Groups.

13. Enter Internal-Zone as the name of the VLAN Group used for Internal network.

14. Choose HANA-Internal, HANA-Storage, MapR-01, MapR-02, NFS-IPMI.

Create VLAN Group ? X

Name :

VLANs

Advanced Filter Export Print No Native VLAN

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	HANA-Internal	<input type="radio"/>
<input type="checkbox"/>	HANA-Replication	<input type="radio"/>
<input checked="" type="checkbox"/>	HANA-Storage	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>
<input checked="" type="checkbox"/>	MapR-01	<input type="radio"/>
<input checked="" type="checkbox"/>	MapR-02	<input type="radio"/>

Create VLAN

< Prev Next > **Finish** Cancel

15. Click Next.

16. Click Next on Add Uplink Ports.

17. Choose the port-channels created for the Internal Zone.

1 Select VLANs

2 Add Uplink Ports

3 Add Port Channels

Create VLAN Group

Port Channels		
Name	Fabric ID	ID
vPC-3...	B	32
vPC-3...	A	31

Selected Port Channels		
Name	Fabric ID	ID
N3k-Uplink	A	5
N3k-Uplink	B	6

< Prev Next > **Finish** Cancel

18. Click Finish.

19. Choose Create VLAN Groups.

20. Enter Client-Zone as the name of the VLAN Group used for Client network.

21. Choose HANA-AppServer, HANA-DataStore.

22. Click Next.

23. Click Next on Add Uplink Ports.

24. Choose port-channels created for Client Zone.

25. Click Finish.

26. Choose Create VLAN Groups.

27. Enter Backup-Network as the name of the VLAN Group used for Backup Network.

28. Choose HANA-Backup.
29. Click Next.
30. Click Next on Add Uplink Ports, since we will use Port-Channel.
31. Choose the port-channels created for Backup Network.
32. Click Finish.
33. Choose Create VLAN Groups.
34. Enter Replication-Network as the name of the VLAN Group used for Replication Network.
35. Choose HANA-Replication.
36. Click Next.
37. Click Next on Add Uplink Ports.
38. Choose the port-channels created for Replication Network.
39. Click Finish.

LAN / LAN Cloud / VLAN Groups

VLAN Groups		Events				
+ - Advanced Filter Export Print						
Name	Native VLAN	Native VLAN DN	Size	VLAN ID	Poolable DN	
▼ LAN Cloud						
▼ VLAN Group Replication-Network			1			
VLAN HANA-Replication				225	fabric/lan/net-HANA-R...	
▼ VLAN Group Internal-Zone			5			
VLAN HANA-Internal				220	fabric/lan/net-HANA-In...	
VLAN HANA-Storage				110	fabric/lan/net-HANA-St...	
VLAN MapR-01				21	fabric/lan/net-MapR-01	
VLAN MapR-02				22	fabric/lan/net-MapR-02	
VLAN NFS-IPMI				197	fabric/lan/net-NFS-IPMI	
▼ VLAN Group Client-Zone			4			
VLAN DMZ				401	fabric/lan/net-DMZ	
VLAN HANA-AppServer				223	fabric/lan/net-HANA-A...	
VLAN HANA-Client				222	fabric/lan/net-HANA-CL...	
VLAN HANA-DataSource				224	fabric/lan/net-HANA-D...	
▼ VLAN Group Backup-Network			1			



For each VLAN Group a dedicated or shared Ethernet Uplink Port or Port Channel can be selected.

Create vNIC Template

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

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

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 HANA-Internal 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. Select Updating Template as the Template Type.
10. Under VLANs, check the checkboxes for HANA-Internal.
11. Set HANA-Internal as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-A.
14. For Network Control Policy, choose Internal from drop-down list.

Create vNIC Template



Name : HANA-Internal

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
<input type="checkbox"/>		
<input checked="" type="checkbox"/>	HANA-Internal	<input checked="" type="radio"/>
<input type="checkbox"/>	HANA-Replication	<input type="radio"/>

OK

Cancel

Create vNIC Template



SELECT	NAME	INCLUDE VLAN
<input checked="" type="checkbox"/>	HANA-Internal	<input checked="" type="radio"/>
<input type="checkbox"/>	HANA-Replication	<input type="radio"/>
<input type="checkbox"/>	HANA-Storage	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	MapR-01	<input type="radio"/>

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 :

15. Click OK to create the vNIC template.



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. Follow the steps 1-15 above to create vNIC template for each Network Interface.

Create a vNIC Template for Storage Network

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 HANA-Storage 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 HANA-Storage.
11. Set HANA-Storage as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-B.
14. For Network Control Policy, choose Internal from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for AppServer Network

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 HANA-AppServer 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 HANA-AppServer.
11. Set HANA-AppServer as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-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

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 HANA-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 HANA-Backup.
11. Set HANA-Backup as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-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 Client Network

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 HANA-Client 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 HANA-Client.
11. Set HANA-Client as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-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

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 HANA-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 HANA-DataSource.
11. Set HANA-DataSource as the native VLAN.

12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-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 Replication Network

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 HANA-Replication 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 HANA-Replication.
11. Set HANA-Replication as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-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 Management Network

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 9000.
13. In the MAC Pool list, choose FI-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 IPMI Inband Access Network

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-IPMI 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-IPMI.
11. Set NFS-IPMI as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-A.

14. For Network Control Policy, choose Internal from drop-down list.
15. Click OK to create the vNIC template.

Create a vNIC Template for MapR-01 Internal Network

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 MapR-01 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 MapR-01.
11. Set MapR-01 as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-A.
14. For Network Control Policy, choose Internal from drop-down list.
15. Click OK to create the vNIC template

Create a vNIC template for MapR-02 Internal Network.

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 MapR-02 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 MapR-02.
11. Set MapR-02 as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, choose FI-B.
14. For Network Control Policy, choose Internal from drop-down list.
15. Click OK to create the vNIC template.

Create Boot Policies

To create Local boot policies, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root > Sub-Organization > HANA.
3. Right-click Boot Policies.
4. Choose Create Boot Policy.
5. Enter Local-Boot as the name of the boot policy.
6. (Optional) Enter a description for the boot policy.
7. Expand the Local Devices drop-down list and choose Add CD/DVD.
8. Expand the Local Devices drop-down list and choose Add Local Disk.
9. Click OK to save the boot policy.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.

The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.

If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.

If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

Add Local Disk

- Add Local LUN
- Add Local JBOD
- Add SD Card
- Add Internal USB
- Add External USB
- Add Embedded Local LUN
- Add Embedded Local Disk

Add CD/DVD

- Add Local CD/DVD
- Add Remote CD/DVD

Boot Order

+ - Advanced Filter Export Print

Name	Order▲	vNIC/v...	Type	WWN	LUN N...	Slot N...	Boot N...	Boot P...	Descri...
CD/DVD	1								
Local Disk	2								

Move Up Move Down Delete

Set Uefi Boot Parameters

Create IP Pool for Inband Management

For SAP HANA High Availability configuration, you will use the IPMI tool, create a block of IP addresses for SAP HANA servers. In the Cisco UCS environment, complete the following steps:



This block of IP addresses should be in the same subnet as the HANA-Admin IP addresses of the HANA Servers.

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose Pools > root > > Sub-Organization > HANA > IP Pools Right-click Create a IP Pool.
3. For the Name enter HANA_IPMI.
4. For Assignment order choose Sequential and click Next.
5. Click Add.
6. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

7. Click OK to create the IP block.
8. Click OK in the confirmation message.

Configure the Inband Profile

To allocate the previously configured IPv4 Address Pool, VLAN, and VLAN Group to the global Inband Profile, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Choose LAN Cloud, on the right pane click the Global Policies tab.
3. On the Global Policies page, under the Inband Profile section:
4. For Inband VLAN Group choose Internal-Zone from the drop-down list.
5. For Network choose NFS-IPMI from the drop-down list.
6. For IP Pool Name choose HANA-IPMI from the drop-down list.
7. Click Save Changes.

LAN / LAN Cloud

LAN Uplinks	VLANs	Server Links	MAC Identity Assignment	IP Identity Assignment	QoS	Global Policies	Faults	Events	FSM
Actions									
Use Global for InBand Profile									
Properties									
MAC Address Table Aging									
Aging Time : <input type="radio"/> Never <input checked="" type="radio"/> Mode Default <input type="radio"/> other									
VLAN Port Count									
VLAN Port Count Optimization : <input type="radio"/> Enabled <input checked="" type="radio"/> Disabled									
Org Permissions									
Org Permissions : <input type="radio"/> Enabled <input checked="" type="radio"/> Disabled									
Inband Profile									
Owner : Local									
Inband VLAN Group : <input type="text" value="Internal-Zone"/>									
Network : <input type="text" value="NFS-IPMI"/>									
IP Pool Name : <input type="text" value="HANA_IPMI"/>									

Create Service Profile Templates for SAP HANA Scale-Out Servers

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA.
3. Right-click HANA.
4. Choose Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter HANA-Server as the name of the service profile template.
 - b. Choose the Updating Template option.
 - c. Under UUID, choose HANA-UUID as the UUID pool.
 - d. Click Next.

Create Service Profile Template ? X

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

Name :

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

The template will be created in the following organization. Its name must be unique within this organization.
Type : Initial Template Updating Template

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

UUID Assignment:

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

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

6. Configure the Storage Provisioning:

- a. Click the Local Disk Configuration Policy tab.
- b. Choose RAID1 for Local Storage field from the drop-down list.
- c. Click Next.

Create Service Profile Template

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

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

Local Storage: **RAID1**

[Create Local Disk Configuration Policy](#)

Mode : **RAID 1 Mirrored**
 Protect Configuration : **Yes**

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

FlexFlash

FlexFlash State : **Disable**

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

FlexFlash RAID Reporting State : **Disable**

< Prev Next > **Finish** Cancel

7. Configure the Networking options:

- a. Keep the default settings for Dynamic vNIC Connection Policy.
- b. Choose the Expert option for How would you like to configure LAN connectivity.
- c. Click the Add button to add a vNIC to the template.
- d. In the Create vNIC dialog box, enter HANA-Internal as the name of the vNIC.
- e. Check the Use vNIC Template checkbox.
- f. In the vNIC Template list, choose HANA-Internal.
- g. In the Adapter Policy list, choose Linux.
- h. Click OK to add this vNIC to the template.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

- i. Repeat the above steps c-h for each vNIC.
8. Add vNIC for HANA-Storage.
9. Add vNIC for HANA-Client.
10. Add vNIC for HANA-AppServer.
11. Add vNIC for HANA-DataSource.
12. Add vNIC for HANA-Replication.
13. Add vNIC for HANA-Backup.
14. Add vNIC for Mgmt.
15. Add vNIC for NFS-IPMI.
16. Review the table in the Networking page to make sure that all vNICs were created.

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

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

Name	MAC Address	Fabric ID	Native VLAN
vNIC HANA-Storage	Derived	derived	
vNIC HANA-Replication	Derived	derived	
vNIC HANA-DataSource	Derived	derived	
vNIC HANA-Client	Derived	derived	
vNIC HANA-Backup	Derived	derived	

[Delete](#) [Add](#) [Modify](#)

[+ iSCSI vNICs](#)

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

17. Click Next.

18. Configure the SAN Connectivity, choose the No vHBAs option for the “How would you like to configure SAN connectivity?” field.

19. Click Next.

20. Set no Zoning options and click Next.

21. Set the vNIC/vHBA placement options:

- a. In the Select Placement list, choose the Specify Manually.
- b. Choose vCon1 and assign the vNICs to the virtual network interfaces policy in the following order:
 - i. HANA-Storage
 - ii. HANA-AppServer
 - iii. HANA-Backup
 - iv. Mgmt

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
HANA-Client
HANA-DataSource
HANA-Internal
HANA-Replication
NFS-IPMI

>> assign >>
<< remove <<

Specific Virtual Network Interfaces (click on a cell to edit)

Name	Order ▲	Selecti...	Admin ...
▼ vCon 1			
		All	
vNIC HANA-Storage	1	ANY	
vNIC HANA-AppServer	2	ANY	
vNIC HANA-Backup	3	ANY	
vNIC Mgmt	4	ANY	
vCon 2			
		All	
		↑ Move Up	↓ Move Down

- c. Choose vCon2 and assign the vNICs to the virtual network interfaces policy in the following order:
- i. HANA-Internal
 - ii. HANA-Client
 - iii. HANA-DataSource
 - iv. HANA-Replication
 - v. NFS-IPMI

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 <<

Specific Virtual Network Interfaces (click on a cell to edit)

Name	Order ▲	Selecti...	Admin ...
▼ vCon 2			
vNIC HANA-Internal	1	ANY	
vNIC HANA-Client	2	ANY	
vNIC HANA-DataSource	3	ANY	
vNIC HANA-Replication	4	ANY	
vNIC NFS-IPMI	5	ANY	

↑ Move Up
 ↓ Move Down

d. Review the table to verify that all vNICs are assigned to the policy in the appropriate order.

e. Click Next.

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

23. Set the server boot order:

a. Choose Local-Boot for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Local-Boot** [Create Boot Policy](#)

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

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

Boot Order

Name	Order	vNIC/vHBA...	Type	WWN	LUN Name	Slot Num...	Boot Name	Boot Path	Description
CD/DVD	1								
Local ...	2								

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

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

b. Click Next.

24. For Maintenance policy:

a. Choose the default Maintenance Policy.

b. Click Next.

25. Specify the server assignment:

a. Choose Up as the power state to be applied when the profile is associated with the server.

b. Expand Firmware Management at the bottom of the page and choose HANA-FW from the Host Firmware list.

c. Click Next.

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: [Create Server Pool](#)

Select the power state to be applied when this profile is associated with the server.
 Up Down

The service profile template is not automatically associated with a server. Either select a server from the list or associate the service profile manually later.

⊖ Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: [Create Host Firmware Package](#)

< Prev Next > **Finish** Cancel

26. For Operational Policies:

- a. In the BIOS Policy list, choose HANA-BIOS.
- b. Expand the External IPMI Management Configuration and choose HANA-IPMI in the IPMI Access Profile.
- c. Choose SoL-Console in the SoL Configuration Profile.
- d. Expand Management IP Address,
- e. In the Outband IPv4 tab choose ext-mgmt in the Management IP Address Policy.
- f. In the Inband IPv4 tab choose NFS-IPMI for Network and choose HANA-IPMI for Management IP Address Policy.

Management IP Address

Outband IPv4 **Inband**

Network : NFS-IPMI ▼

Inband IPv4 Inband IPv6

Management IP Address Policy: HANA_IPMI(14/48) ▼

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

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

[Create IP Pool](#)

27. Expand Power Control Policy Configuration and choose No-Power-Cap in the Power Control Policy list.

Create Service Profile Template ? X

Optionally specify information that affects how the system operates.

BIOS Configuration

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

BIOS Policy :

External IPMI Management Configuration

If you want to access the CIMC on the server externally, select an IPMI access profile. The users and passwords in that profile will be populated into the CIMC when the profile is associated with the server.

IPMI Access Profile : [Create IPMI Access Profile](#)

To enable Serial over LAN access to the server, select an SoL configuration profile.

SoL Configuration Profile:

[Create Serial over LAN Policy](#)

Name : **SoL**

Description :

Management IP Address

Monitoring Configuration (Thresholds)

28. Click Finish to create the service profile template.

29. Click OK in the confirmation message.

If you are using C480-M5 as HANA Servers, create the service profile template and modify the vNIC placement policy.

To create a clone service profile template created, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA.
3. Right-click Service Template HANA-Server.
4. Choose Create a Clone.
5. Enter HANA-C480 for Clone name and Choose root as Org.
6. Click OK to Create the clone of HANA-C480.
7. Click OK to confirm.
8. On the cloned Service Profile Template, choose HANA-C480 in the Navigation pane and click Network tab.

9. Under Actions Click Modify vNIC/vHBA Placement.
10. Choose vCon1 and assign the vNICs to the virtual network interfaces policy in the following order:
 - a. HANA-Internal
 - b. HANA-Storage
 - c. HANA-AppServer
 - d. HANA-Client
 - e. Mgmt
 - f. HANA-Backup
 - g. HANA-DataSource
 - h. HANA-Replication
 - i. NFS-IPMI
11. Click OK to complete the vNIC/vHBA Placement policy.
12. Click OK to confirm.

Create Service Profile from the Template

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA > Service Template HANA-Server.
3. Right-click Service Template HANA-Server and choose Create Service Profiles from Template.
4. Enter HANA-Server0 as the service profile prefix.
5. Enter 1 as Name Suffix Starting Number.
6. Enter 1 as the Number of Instances.
7. Click OK to create the service profile.

Create Service Profiles From Template ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

To create service profiles from the service profile template HANA-C480, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA > Service Template HANA-C480.
3. Right-click Service Template HANA-C480 and choose Create Service Profiles from Template.
4. Enter HANA-Server0 as the service profile prefix.
5. Enter 2 as Name Suffix Starting Number.
6. Enter 1 as the Number of Instances.
7. Click OK to create the service profile.

To associate service profile created for a specific slot, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile > root > Sub-Organization > HANA > HANA-Server01.
3. Right-click HANA-Server01 and choose Change Service Profile Association.
4. For Server Assignment Choose, choose existing Server for the drop-down list.

5. Click All Servers.
6. Choose the Server as recommended.
7. Repeat for above steps 1-6 for each HANA Servers.

Create Service Profile Templates for MapR Servers

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA.
3. Right-click HANA.
4. Choose Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter MapR as the name of the service profile template.
 - b. Choose the Updating Template option.
 - c. Under UUID, choose HANA-UUID as the UUID pool.
 - d. Click Next.
6. Configure the Storage Provisioning:
 - a. Click the Local Disk Configuration Policy tab.
 - b. Choose MapR for Local Storage field from the dropdown menu.
 - c. Click Next.

Create Service Profile Template

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

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

Local Storage:

[Create Local Disk Configuration Policy](#)

Mode : **Any Configuration**
 Protect Configuration : **Yes**

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

FlexFlash

FlexFlash State : **Disable**

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

FlexFlash RAID Reporting State : **Disable**

< Prev Next > **Finish** Cancel

7. Configure the Networking options:

- a. Keep the default settings for Dynamic vNIC Connection Policy.
- b. Choose the Expert option for How would you like to configure LAN connectivity.
- c. Click the Add button to add a vNIC to the template.
- d. In the Create vNIC dialog box, enter MapR-01 as the name of the vNIC.
- e. Check the Use vNIC Template checkbox.
- f. In the vNIC Template list, choose MapR-01
- g. In the Adapter Policy list, choose Linux.
- h. Click OK to add this vNIC to the template.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

- i. Repeat the above steps c-h for each vNIC.
- 8. Add vNIC for MapR-02.
- 9. Add vNIC for HANA-Storage.
- 10. Add vNIC for Mgmt.
- 11. Review the table in the Networking page to make sure that all vNICs were created.

Create Service Profile Template ? X

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

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

Name	MAC Address	Fabric ID	Native VLAN
vNIC Mgmt	Derived	derived	
vNIC HANA-Storage	Derived	derived	
vNIC MapR-02	Derived	derived	
vNIC MapR-01	Derived	derived	

Delete + Add Modify

+ iSCSI vNICs

12. Click Next.

13. Configure the SAN Connectivity, choose the No vHBAs option for the “How would you like to configure SAN connectivity?” field.

14. Click Next.

15. Set no Zoning options and click Next.

16. Set the vNIC/vHBA placement options, keep the default values.

17. Rearrange the order, using Move Up and Move Down icons.

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](#)

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vNIC MapR-01	Derived	1
vNIC MapR-02	Derived	2
vNIC HANA-Storage	Derived	3
vNIC Mgmt	Derived	4

↑ Move Up
↓ Move Down
🗑 Delete
🔄 Reorder
ⓘ Modify

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

18. Click Next.

19. Keep the default values on the vMedia Policy, click Next

20. To set the server boot order, choose Local-Boot for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

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

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

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vHB...	Type	WWN	LUN Name	Slot Num...	Boot Name	Boot Path	Description
CD/DVD	1								
Local ...	2								

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

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

21. Click Next.

22. For Maintenance policy: Choose the default Maintenance Policy.

23. Click Next.

24. Specify the server assignment: Choose Up as the power state to be applied when the profile is associated with the server.

25. Expand Firmware Management at the bottom of the page and choose HANA-FW from the Host Firmware list.

26. Click Next.

27. For Operational policies, In the BIOS Policy list, choose MapR-BIOS.

28. Expand Management IP Address, in the Outband IPv4 tab choose ext-mgmt in the Management IP Address Policy.

29. Expand Power Control Policy Configuration and choose No-Power-Cap in the Power Control Policy list.

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 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

< Prev Next > **Finish** Cancel

30. Click Finish to create the service profile template.

31. Click OK in the confirmation message.

Create Service Profiles

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile Templates > root > Sub-Organization > HANA > Service Template MapR.
3. Right-click MapR and choose Create Service Profiles from Template.
4. Enter appropriate name for the service profile prefix. For example, MapR-0.
5. Enter 1 as Name Suffix Starting Number.
6. Enter appropriate number of service profile to be created in the Number of Instances. For example, 4.
7. Click OK to create the service profile.

Create Service Profiles From Template ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

OK

Cancel

To associate service profile created for a specific slot, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile > root > Sub-Organization > HANA > MapR-01.
3. Right-click MapR-01 and choose Change Service Profile Association.
4. For Server Assignment, choose the existing Server from the drop-down list.
5. Click All Servers.
6. Choose the C240-M5 Rack Mount Server Rack ID 1.
7. Repeat the above steps 1-6 for each MapR Server.

MapR Storage Configuration

The section describes the procedure for Installing MapR Data Platform on Cisco UCS C240-M5 Servers. Each server has 2 x Internal SSD Boot drives where the Operating System is installed with RAID 1. For MapR Storage 24 x 1.8 TB 10k rpm 4K disks are configured in 4 x RAID 5 group with 6 disks.

MapR Server RAID Configuration

To configure RAID on the Cisco UCS C240 Servers used for MapR Data Platform, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile > root > Sub-Organization > HANA > MapR-01.
3. Click KVM Console.
4. When the KVM Console is launched, click Boot Server.

```
Cisco Systems, Inc.  
Configuring and testing memory..  
Configuring platform hardware...
```

UCSM

5. When prompted Press <Ctrl><R> to Run MegaRAID Configuration Utility.

```

ID   LUN  VENDOR  PRODUCT                REVISION  CAPACITY
---  ---  ---     ---                    -
49   0    HGST    HUC101818CS4200      A703      1716957MB
50   0    ATA     INTEL SSDSC2BB15     CS01      143089MB
51   0    HGST    HUC101818CS4200      A703      1716957MB

26 JBOD(s) found on the host adapter
2 JBOD(s) handled by BIOS

0 Virtual Drive(s) found on the host adapter.
0 Virtual Drive(s) handled by BIOS

Press <Ctrl><R> to Run MegaRAID Configuration Utility

```

- In VD Mgmt Tab, highlight Cisco 12G Modular Raid Con and press F2 and choose Make Unconfigured Good.

```


Cisco 12G Modular Raid Con BIOS Configuration Utility 7.01-0801
UD Mgmt  PD Mgmt  Ctrl Mgmt  Properties
-----
Virtual Drive Management
[-] Cisco 12G Modular Raid Con(Bus 0x18, Dev 0x00)
  └─ No Configuration Present !

Create Virtual Drive
Clear Configuration
Foreign Config
Manage Preserved Cache
Drive Security
Disable Data Protection
Make Unconfigured Good
Make JBOD
Advanced Software Options


F1-Help F2-Operations F5-Refresh Ctrl-N-Next Page Ctrl-P-Prev Page F12-Ctrl

```


7. In the Convert JBODs to Unconfigured Good choose All JBOD Drives and enter OK.




8. In the VD Mgmt press F2 for Operations and choose Create Virtual Drive.



9. Choose the following option to create RAID 1 virtual drive with 2 disks for Boot:
 - a. For RAID Level Choose RAID-1.
 - b. Choose the last two drives in Slot 25 and 26.
 - c. Give a Name for virtual drive 0.




10. Choose Advanced option:
 - a. For Strip Size choose 1MB.
 - b. For Read Policy Choose Ahead.
 - c. For Write Policy choose Write Back with BBU.
 - d. For I/O Policy choose Direct.
 - e. Keep the Default Disk cache Policy and Emulation.
 - f. Choose Initialize.
 - g. Click OK.



11. In the VD Mgmt press F2 for Operations and choose Create Virtual Drive.


12. Choose the following option to create RAID 5 virtual drive with 6 disks:

- a. For RAID Level Choose RAID-5.
- b. Choose the first 6 disks.
- c. Give a Name for virtual drive 1.



13. Choose Advanced option:

- a. For Strip Size choose 1MB.
- b. For Read Policy Choose Ahead.
- c. For Write Policy choose Write Back with BBU.
- d. For I/O Policy choose Direct.
- e. Keep the Default Disk cache Policy and Emulation.
- f. Choose Initialize.
- g. Click OK.



14. To create the second RAID 5 Virtual Drive press F2 for Operations and choose Create Virtual Drive.

15. Create a RAID 5 virtual drive with 6 disks:

- a. For RAID Level Choose RAID-5.
- b. Choose the next 6 disks.
- c. Give a Name for virtual drive 2.
- d. Choose Advanced option.
- e. For Strip Size choose 1MB.
- f. For Read Policy Choose Ahead.
- g. For Write Policy choose Write Back with BBU.
- h. For I/O Policy choose Direct.
- i. Keep the Default Disk cache Policy and Emulation.
- j. Choose Initialize.
- k. Click OK.

16. To create the third RAID 5 Virtual Drive press F2 for Operations and choose Create Virtual Drive.

17. Create a RAID 5 virtual drive with 6 disks:

- a. For RAID Level Choose RAID-5.

- b. Choose the next 6 disks.
 - c. Give a Name for virtual drive 3.
 - d. Choose Advanced option.
 - e. For Strip Size choose 1MB.
 - f. For Read Policy Choose Ahead.
 - g. For Write Policy choose Write Back with BBU.
 - h. For I/O Policy choose Direct.
 - i. Keep the Default Disk cache Policy and Emulation.
 - j. Choose Initialize.
 - k. Click OK.
18. To create the fourth RAID 5 Virtual Drive press F2 for Operations and choose Create Virtual Drive.
19. Create a RAID 5 virtual drive with 6 disks:
- a. For RAID Level Choose RAID-5.
 - b. Choose the last 6 disks.
 - c. Give a Name for virtual drive 4.
 - d. Choose Advanced option.
 - e. For Strip Size choose 1MB.
 - f. For Read Policy Choose Ahead.
 - g. For Write Policy choose Write Back with BBU.
 - h. For I/O Policy choose Direct.
 - i. Keep the Default Disk cache Policy and Emulation.
 - j. Choose Initialize.
 - k. Click OK.
20. Press esc to exist MegaRAID Configuration Utility.

```

Cisco 12G Modular Raid Con BIOS Configuration Utility 7.01-0801
UD Mgmt PD Mgmt Ctrl Mgmt Properties
Virtual Drive Management
[-] Cisco 12G Modular Raid Con(Bus 0x18, Dev 0x00)
  [-] Drive Group: 0, RAID 1
    [-] Virtual Drives
      ID: 0, VD0, 138.766 GB
    [+] Drives
    [+] Available size: 0.000 KB
    Hot spare drives
  [+] Drive Group: 1, RAID 5
  [+] Drive Group: 2, RAID 5
  [+] Drive Group: 3, RAID 5
  [-] Drive Group: 4, RAID 5
    [-] Virtual Drives
      ID: 4, VD4, 8.180 TB
    [+] Drives
    [+] Available size: 35.000 MB
    Hot spare drives

Drive Group 0:
Virtual Drives: 1
Drives: 2
Free Cap.: 0.000 KB
Free Areas: 0
Secured : N/A
Protection : N/A

F1-Help F2-Operations F5-Refresh Ctrl-N-Next Page Ctrl-P-Prev Page F12-Ctrl

```

21. Press Ctrl+Alt+Del to reboot the server.

22. Repeat the steps 1-21 for each MapR Servers.


MapR Server Operating System Installation

To install RedHat Enterprise Linux 7.4 Operation System on Internal Boot Drives, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile > root > Sub-Organization > HANA > MapR-01.
3. Click KVM Console.
4. When the KVM Console is launched, click Boot Server.
5. Click Virtual Media > Activate Virtual Devices:
 - a. Choose the option Accept this Session for Unencrypted Virtual Media Session and then click Apply.
 - b. Click Virtual Media and Choose Map CD/DVD.
 - c. Click Browse to navigate ISO media location.
 - d. Click Map Device.
6. On the Initial screen choose Install to begin the installation process.

7. Choose Language and click Continue.

The Installation Summary page displays and this is where you will configure various features.



8. Click Date & Time; choose the appropriate timezone and click Done.

9. Click Keyboard; choose Keyboard layout and click Done.

10. Under Software Menu, click Software selection.

11. In the Base Environment choose Infrastructure Server.

12. For Add-Ons for Selected Environment choose Large Systems Performance, Network File System Client, Performance Tools, Compatibility Libraries and click Done.

SOFTWARE SELECTION RED HAT ENTERPRISE LINUX 7.4 INSTALLATION

Done us Help

Base Environment

- Minimal Install**
Basic functionality.
- Infrastructure Server**
Server for operating network infrastructure services.
- File and Print Server**
File, print, and storage server for enterprises.
- Basic Web Server**
Server for serving static and dynamic internet content.
- Virtualization Host**
Minimal virtualization host.
- Server with GUI**
Server for operating network infrastructure services, with a GUI.

Add-Ons for Selected Environment

- RDMA-based InfiniBand and iWARP fabrics.
- Java Platform**
Java support for the Red Hat Enterprise Linux Server and Desktop Platforms.
- Large Systems Performance**
Performance support tools for large systems.
- Load Balancer**
Load balancing support for network traffic.
- MariaDB Database Server**
The MariaDB SQL database server, and associated packages.
- Network File System Client**
Enables the system to attach to network storage.
- Performance Tools**
Tools for diagnosing system and application-level performance problems.
- PostgreSQL Database Server**
The PostgreSQL SQL database server, and associated packages.
- Print Server**
Allows the system to act as a print server.
- Remote Management for Linux**
Remote management interface for Red Hat Enterprise Linux, including OpenLMI and SNMP.
- Virtualization Hypervisor**
Smallest possible virtualization host installation.
- Compatibility Libraries**
Compatibility libraries for applications built on previous versions of Red Hat Enterprise Linux.
- Development Tools**
A basic development environment.
- Security Tools**
Security tools for integrity and trust verification.
- Smart Card Support**
Support for using smart card authentication.

13. Under System, click Installation destination.

INSTALLATION DESTINATION RED HAT ENTERPRISE LINUX 7.4 INSTALLATION

[Done](#)  [Help!](#)

Device Selection

Select the device(s) you'd like to install to. They will be left untouched until you click on the main menu's "Begin Installation" button.

Local Standard Disks

138.77 GiB	8377.19 GiB	8377.19 GiB	8377.19 GiB	Cisco UCSC-RAID12GP-4G
				
Cisco UCSC-RAID12GP-4G sda / 36.58 GiB free	Cisco UCSC-RAID12GP-4G sdb / 8377.19 GiB free	Cisco UCSC-RAID12GP-4G sdc / 8377.19 GiB free	Cisco UCSC-RAID12GP-4G sdd / 8377.19 GiB free	Cisco UCSC-RAID12GP-4G sde / 8377.19 GiB free

Disks left unselected here will not be touched.

Specialized & Network Disks

[Add a disk...](#)

Disks left unselected here will not be touched.

Other Storage Options

Partitioning

Automatically configure partitioning. I will configure partitioning.

I would like to make additional space available.

Encryption

Encrypt my data. You'll set a passphrase next.

[Full disk summary and boot loader...](#) 1 disk selected; 138.77 GiB capacity; 36.58 GiB free [Refresh...](#)


14. Choose the Boot drives with RAID 1 created.

15. From the Other Storage Options choose 'I will configure partitioning' and click Done.


16. In the Manual Partitioning Screen, choose Standard Partition for New mount points will use the following partitioning scheme.

The screenshot shows the 'MANUAL PARTITIONING' screen for 'RED HAT ENTERPRISE LINUX 7.4 INSTALLATION'. At the top left is a 'Done' button, and at the top right is a 'Help!' button. The main content area is titled 'New Red Hat Enterprise Linux 7.4 Installation' and contains the following text: 'You haven't created any mount points for your Red Hat Enterprise Linux 7.4 installation yet. You can:'. Below this are two bullet points: '• [Click here to create them automatically.](#)' and '• Create new mount points by clicking the '+' button.' The text continues: 'New mount points will use the following partitioning scheme:'. Below this is a dropdown menu currently set to 'Standard Partition'. At the bottom of the main content area are three buttons: '+', '-', and a refresh icon. Below these buttons are two boxes: 'AVAILABLE SPACE 138.77 GiB' and 'TOTAL SPACE 138.77 GiB'. At the bottom left, it says '1 storage device selected'. At the bottom right, there is a 'Reset All' button. A note on the right side of the screen reads: 'When you create mount points for your Red Hat Enterprise Linux 7.4 installation, you'll be able to view their details here.'

17. Click the + symbol to add a new partition.
18. Choose the mount point /boot.
19. Enter the Desired capacity as 200 MB and click Add Mount Point.



20. Choose the filesystem ext3.
21. Click the + symbol to add a new partition.
22. Choose the mount point as swap.
23. Enter the Desired capacity as 2048 MB and click Add Mount Point.
24. Choose the filesystem as swap.
25. Click the + symbol to add / (root) partition.
26. Choose the mount point as /.
27. Enter the Desired capacity as 102400 MiB and click Add Mount Point.
28. Choose the filesystem ext3.




29. Click Done on the top left corner of the screen.
30. Review the partition layout and the size.
31. Click Accept Changes to proceed with the next steps.
32. Click KDUMP from the below screen.
33. Deselect Enable kdump.
34. Click Security policy, choose Apply Security policy to OFF and click Done.
35. Click Done on the top left corner of the screen.
36. Click Network & Hostname.
37. Enter the Host name and click Apply.

 Starting Red Hat Enterprise Linux 7.x, the system assigns the Ethernet device names based on the bus where they are present. The default Ethernet device naming convention like eth0, eth1 ...ethx is deprecated. This can be addressed once the OS is installed and active.

38. IP address will be assigned once the OS is installed. Click Done.


39. Click Done on the top left corner of the screen.

40. Review the installation summary and click Begin Installation.




41. The next screen shows the start of the OS installation.

42. Click Root Password.




43. Enter the Root Password and Confirm.

44. Click Done on the top left corner of the window.



45. The installation will start and continue.

46. When the installation is complete click Reboot to finish the installation.



Operating System Network Configuration

To customize the server in preparation for the MapR Installation, complete the following steps:

Hostnames

 The operating system must be configured in such a way that the short name of the server is displayed for the command **hostname -s**, and the fully qualified host name is displayed with the command **hostname -f**

1. Use the KVM console to log in to the installed system as the user root and the password <<var_sys_root-pw>>.
2. Set the hostname using hostnamectl:

```
hostnamectl set-hostname <<hostname>>
```

```
hostnamectl set-hostname <<hostname>>
```

3. Edit the Hostname:

```
vi /etc/HOSTNAME
<<hostname>>.<<Domain Name>>
```

Configure the Network

Each MapR Server is configured with 4 vNIC device. Table 11 lists the IP Address information required to configure the IP address on the Operating System. IP Address and Subnet Mask provided below are for example only; please configure IP address as per your environment.

Table 11 List the IP Address for SAP HANA Server

vNIC Name	VLAN ID	IP Address Range	Subnet Mask
MapR-01	<<var_mapr-01_vlan_id>>	10.21.21.101 to 10.21.21.104	255.255.255.0
MapR-02	<<var_mapr-02_vlan_id>>	10.22.22.101 to 10.22.22.104	255.255.255.0
HANA-Storage	<<var_storage_vlan_id>>	192.168.110.101 to 192.168.110.104	255.255.255.0
Management	<<var_mgmt_vlan_id>>	192.168.196.101 to 192.168.196.104	255.255.0.0

1. In RHEL 7, systemd and udev support a number of different naming schemes. By default, fixed names **are assigned based on firmware, topology, and location information, like 'enp72s0'**. With this naming convention, though names stay fixed even if hardware is added or removed it is often harder to read unlike traditional kernel-native ethX naming "eth0". **Another way to name network interfaces, "biosdev-names", is already available with installation.**
2. Configure boot parameters "net.ifnames=0 biosdevname=0" to disable both, to get the original kernel native network names.
3. Also, IPV6 support could be disabled at this time as we use IPV4 in the solution. This can be done by appending ipv6.disable=1 to GRUB_CMDLINE_LINUX as shown below:

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

4. To Run the grub2-mkconfig command to regenerate the grub.cfg file:

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

5. Reboot system to effect the changes.
6. To configure the network interface on the OS, it is required to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS.

7. From the OS, execute the following command to get list of Ethernet device with MAC Address:

```
[root@mapr31 ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:2d brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:24 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:25 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:2e brd ff:ff:ff:ff:ff:ff
```

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

9. Expand Servers > Service Profile > root > Sub-Organization > HANA > MapR-01.

10. Click + to Expand. Click vNICs.

11. On the right pane list of the vNICs with MAC Address are listed.

vNICs

▼ Advanced Filter ↑ Export 🖨 Print				
Name	MAC Address	Desired O... ▲	Actual Order	Fabric ID
vNIC MapR-01	00:25:B5:7A:00:2D	1	1	A B
vNIC MapR-02	00:25:B5:7B:00:24	2	2	B A
vNIC HANA-Storage	00:25:B5:7B:00:25	3	3	B A
vNIC Mgmt	00:25:B5:7A:00:2E	4	4	A B

12. Note the MAC Address of the MapR-01 vNIC is “00:25:B5:7A:00:2D”.

13. By comparing the MAC Address on the OS and Cisco UCS, eth0 on OS will carry the VLAN for MapR-01.

14. Go to the network configuration directory and create a configuration for eth0:

```
cd /etc/sysconfig/network-scripts/
vi ifcfg-eth0
##
# MapR-01 Network
##
DEVICE=eth0
TYPE=Ethernet
ONBOOT=yes
BOOTPROTO=static
IPV6INIT=no
USERCTL=no
NM_CONTROLLED=no
```

```
IPADDR=<<IP address for MapR-01 example:10.21.21.101/24>>
NETMASK=<<subnet mask for MapR-01 255.255.255.0>>
```

15. Repeat the steps 11 to 18 for each vNIC interface.

16. Add default gateway:

```
vi /etc/sysconfig/network

NETWORKING=yes
HOSTNAME=<<HOSTNAME>>
GATEWAY=<<IP Address of default gateway>>
```

DNS

Domain Name Service configuration must be done based on the local requirements.

Configuration Example

Add DNS IP if it is required to access internet:

```
vi /etc/resolv.conf

DNS1=<<IP of DNS Server1>>
DNS2=<<IP of DNS Server2>>
DOMAIN= <<Domain_name>>
```

All nodes should be able to resolve internal network IP address, below is an example of 4 node host file with the entire network defined in the /etc/hosts file.

Example:

```
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
#
## MapR 01
#
10.21.21.101 mapr31.ciscolab.local mapr31
10.21.21.102 mapr32.ciscolab.local mapr32
10.21.21.103 mapr33.ciscolab.local mapr33
10.21.21.104 mapr34.ciscolab.local mapr34
#
## MapR 02
#
10.22.22.101 mapr31b.ciscolab.local mapr31b
10.22.22.102 mapr32b.ciscolab.local mapr32b
10.22.22.103 mapr33b.ciscolab.local mapr33b
10.22.22.104 mapr34b.ciscolab.local mapr34b
#
## HANA Storage Network
#
192.168.110.211 cishanaso11s.ciscolab.local cishanaso11s
192.168.110.212 cishanaso12s.ciscolab.local cishanaso12s
192.168.110.213 cishanaso13s.ciscolab.local cishanaso13s
192.168.110.214 cishanaso14s.ciscolab.local cishanaso14s
192.168.110.215 cishanaso15s.ciscolab.local cishanaso15s
192.168.110.216 cishanaso16s.ciscolab.local cishanaso16s
192.168.110.217 cishanaso17s.ciscolab.local cishanaso17s
192.168.110.218 cishanaso18s.ciscolab.local cishanaso18s
#
## MapR Storage
#
192.168.110.101 mapr31s.ciscolab.local mapr31s
192.168.110.102 mapr32s.ciscolab.local mapr32s
```

```

192.168.110.103 mapr33s.ciscolab.local mapr33s
192.168.110.104 mapr34s.ciscolab.local mapr34s
#
## MapR Storage Virtual IP
#
192.168.110.71 maprvip31
192.168.110.72 maprvip32
192.168.110.73 maprvip33
192.168.110.74 maprvip34
192.168.110.75 maprvip35
192.168.110.76 maprvip36
192.168.110.77 maprvip37
192.168.110.78 maprvip38
#
## Management
192.168.73.101 mapr31m.ciscolab.local mapr31m
192.168.73.102 mapr32m.ciscolab.local mapr32m
192.168.73.103 mapr33m.ciscolab.local mapr33m
192.168.73.104 mapr34m.ciscolab.local mapr34m

```

RedHat System Update and OS Customization for MapR Servers

To update and customize the RedHat System for MapR Servers, complete the following steps:

1. To patch the system, you must first update the repository.



The installed system does not include any update information. To patch the Red Hat system, it must be registered and attached to a valid subscription. The following code will register the installation and update the repository information:

2. Register all nodes to get access to the standard RHEL channels:

```

subscription-manager register --auto-attach
Username: <<username>>
Password: <<password>>

```

3. Update only the OS kernel and firmware packages to the latest release that appeared in RHEL 7.4. Set the release version to 7.4:

```
subscription-manager release --set=7.4
```

4. Update the OS:

```
yum -y update
```

5. Install additional groups and packages:

```

yum -y install nfs-utils bash tuned rpcbind dmidecode glibc glibc-common glibc-headers glibc-devel hdparm
initscripts iputils irqbalance libgcc libstdc++ redhat-lsb-core rpm-libs sdparm shadow-utils syslinux unzip
zip nc mttools syslinux-nonlinux nss python-pycurl openssh-clients openssh-server openssl sshpass sudo wget
which java-1.8.0-openjdk java-1.8.0-openjdk-headless java-1.8.0-openjdk-devel ipmitool

```

6. Disable SELinux:

```

sed -i 's/^SELINUX=enforcing/SELINUX=disabled/g' /etc/sysconfig/selinux
sed -i 's/^SELINUX=permissive/SELINUX=disabled/g' /etc/sysconfig/selinux
sed -i 's/^SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
sed -i 's/^SELINUX=permissive/SELINUX=disabled/g' /etc/selinux/config

```

7. Disable transparent hugepages [THP] and configure C-States for lower latency in Linux. Modify the file `/etc/default/grub` and append the following parameter to the line starting with "GRUB_CMDLINE_LINUX":

```
transparent_hugepage=never
```

8. To Run the `grub2-mkconfig` command to regenerate the `grub.cfg` file:

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

9. Add `mapr` user and group. Make sure that `<uid>` and `<gid>` are greater than 1000 and the same on all MapR cluster nodes. The `mapr` defaults of 2000 for both `uid` and `gid` will ensure consistency with other default MapR clusters for inter-cluster mirroring:

```
groupadd mapr -g <gid>
useradd mapr -u <uid> -g <gid>
```

10. Activate `rpcbind`:

```
systemctl start rpcbind
systemctl enable rpcbind
```

11. Append following parameters in `/etc/sysctl.conf`:

```
sunrpc.tcp_slot_table_entries = 128
sunrpc.tcp_max_slot_table_entries = 128
net.ipv4.tcp_slow_start_after_idle = 0
```

12. Enable tuned profile:

```
tuned-adm profile enterprise-storage
```

13. Disable Firewall:

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

14. Reboot the OS by issuing `reboot` command

15. Optional: old kernels can be removed after OS update:

```
package-cleanup --oldkernels --count=1 -y
```

Install Cisco eNIC Driver

To download the Cisco UCS Drivers ISO bundle, which contains most of the Cisco UCS Virtual Interface Card drivers, complete the following steps:

1. In a web browser, navigate to <https://www.cisco.com>.
2. Under Support, click All Downloads.
3. In the product selector, click Products, then click Server - Unified Computing.

4. If prompted, enter your Cisco.com username and password to log in.



You must be signed in to download Cisco Unified Computing System (UCS) drivers.

5. Cisco UCS drivers are available for both Cisco UCS B-Series Blade Server Software and Cisco UCS C-Series Rack-Mount UCS-Managed Server Software.
6. Click UCS B-Series Blade Server Software.
7. Click Cisco Unified Computing System (UCS) Drivers.



The latest release version is selected by default. This document is built on Version 3.2(2a).

8. Click 3.2(2a) Version.
9. Download ISO image of Cisco Unified Computing System (UCS) Drivers.
10. Choose ISO image of UCS-related linux drivers only and click Download and follow the prompts to complete your driver download.
11. After the download is complete browse to ucs-cxxx-drivers-linux.3.1.2b\Network\Cisco\VIC\RHEL\RHEL7.4 and copy kmod-enic-2.3.0.44-rhel7u4.el7.x86_64.rpm to each server.
12. ssh to the Server on the Management IP as root.
13. Update the enic driver with the following command:

```
rpm -Uvh /tmp/kmod-enic-2.3.0.44-rhel7u4.el7.x86_64.rpm
```

14. Update the enic driver on all the MapR Servers.

Network Time

The configuration of NTP is important and must be performed on all systems. To configure network time, complete the following steps:

1. Install NTP-server with utilities.

```
yum -y install ntp ntpdate
```

2. Configure NTP by adding at least one NTP server to the NTP config file /etc/ntp.conf:

```
vi /etc/ntp.conf
server <NTP-SERVER1 IP>
server <NTP-SERVER2 IP>
```

3. Stop the NTP services and update the NTP Servers:

```
systemctl stop ntpd
ntpdate ntp.example.com
```

4. Start NTP service and configure it to be started automatically:

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

SSH Keys

The SSH Keys must be exchanged between all MapR servers for user 'root'. To exchange the SSH keys, complete the following steps:

1. Generate the rsa public key by executing the following command:

```
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:
62:8a:0b:c4:4c:d7:de:8d:dc:e7:b7:17:ca:c7:02:94 root@mapr31.ciscolab.local
The key's randomart image is:
+--[ RSA 2048]-----+
| |
| . |
| . . . . |
|+ . . o + E |
| + .o+So.. |
|. . o . o . . |
|. . . .o.o .|
|. . .+.+ |
|. .+ |
+-----+
```

2. Exchange the rsa public key by executing the following command from First server to rest of the servers. This helps ensure that every host can establish a password-less ssh connection to every other host:

“ssh-copy-id -i /root/.ssh/id_rsa.pub mapr31”

```
ssh-copy-id -i /root/.ssh/id_rsa.pub mapr32
```

3. Repeat the steps above for all the servers in the MapR cluster.

MapR Installation

When installing the MapR software, it requires a repository with the MapR software packages. If the MapR cluster nodes have internet access, it is recommended to prepare an online repository. If the MapR cluster nodes do not have internet access, you can download the packages from a computer which has internet access, unzip and copy them to a server which is reachable by the MapR cluster nodes or to a MapR cluster node itself and prepare an offline repository.



The following information is required to setup the MapR Cluster.

Cluster IDs

Cluster Name: example: mapr600

Virtual IP Range

For High Availability and network bandwidth distribution, virtual IPs are created on the MapR cluster. These virtual IPs are bound to HANA Storage Ethernet device. HANA servers will use these virtual IPs to mount the NFS volumes. If a MapR node is failed, the virtual IPs are moved to other active nodes and the HANA servers will continue to access NFS volumes without any disruption. Its recommended to create two virtual IP address per MapR server, for example if there are 4 MapR nodes in the cluster – Create 8 Virtual IPs.

IP Address: example: 192.168.110.71-78

Subnet Mask: example: 255.255.255.0 or /24

Preparing Online Repository

On every MapR cluster node, create the file `/etc/yum.repos.d/mapr.repo` with the following content. If a proxy server is not necessary, delete the two lines beginning with "proxy=":

```
[maprtech]
name=MapR Technologies
baseurl=http://package.mapr.com/releases/v6.0.0/redhat/
enabled=1
gpgcheck=0
protect=1
proxy=http://proxy:port
```

Verify the repository:

```
yum repolist
```

Installing MapR Packages

Packages will be installed based on the role of the node. To install the MapR packages, complete the following steps:

1. Install on MapR NFS and MapR Fileserver on all nodes in the MapR Cluster:

```
yum -y install mapr-nfs mapr-fileserver
```

2. Install zookeeper service on last three nodes of the MapR Cluster:

```
yum -y install mapr-zookeeper
```

3. Install cldb service on first two nodes of the MapR Cluster:

```
yum -y install mapr-cldb
```

4. Install webserver service on all the nodes or at least on two nodes in the MapR Cluster:

```
yum -y install mapr-webserver
```

Initial Configuration

After all packages are installed, complete the following steps to configure MapR cluster:



WARNING! Any changes after these steps without re-running the configuration script might result in an unstable cluster state.

1. Increase memory usage of MapR instance:

```
sed -i 's/service.command.mfs.heapsize.percent=.*service.command.mfs.heapsize.percent=85/g' /opt/mapr/conf/warden.conf
```

2. Define the MapR subnets to use:

```
sed -i 's%s#export MAPR_SUBNETS=%export MAPR_SUBNETS=<<internal_network_01>>/<<subnet_short>>, <<internal_network_02>>/<<subnet_short>> %g' /opt/mapr/conf/env.sh
```

Example:

```
sed -i -e 's%s#export MAPR_SUBNETS=%export MAPR_SUBNETS=10.21.221.0/24,10.22.222.0/24 %g' /opt/mapr/conf/env.sh
```

3. Modify the DrCache size of the NFS server:

```
sed -i 's/#DrCacheSize =.*DrCacheSize = 614400/g' /opt/mapr/conf/nfsserver.conf
```

4. Create the /mapr directory and the /opt/mapr/conf/mapr_fstab file:

```
mkdir -p /mapr
echo "localhost:/mapr /mapr hard,nolock" > /opt/mapr/conf/mapr_fstab
```

5. Run the configuration script. -C followed by the cldb nodes, -Z followed by the zookeeper nodes, -N gives the cluster name, and -no-autostart hinders the MapR software to start automatically. If you use hostnames instead of IP addresses for -C and -Z, be sure to use hostnames which reflect to the internal cluster network IP addresses:

```
/opt/mapr/server/configure.sh -C <<cldbnode1>>,<<cldbnode2>> -Z <<zookeepernode1>>,<<zookeepernode2>>,<<zookeepernode3>> -N <<clustername>> -noautostart
```

Example:

```
/opt/mapr/server/configure.sh -C mapr31, mapr32 -Z mapr32, mapr33, mapr34 -N mapr600 -no-autostart
```

6. Create a file with a disk list, containing all Virtual Drives in RAID 5 created:

```
lsblk | grep 8.2T | awk '{print "/dev/"$1}' > /tmp/disk.list
```

7. Set up the disks in the disk.list file with a storage pool of 1 disk each on all nodes:



WARNING! The disks will be wiped immediately!

```
/opt/mapr/server/disksetup -W 1 -F /tmp/disk.list
```

8. Repeat Step 1 through 7 on each MapR node in the cluster.

Starting up MapR cluster

To issue the following commands to start the cluster, complete the following steps:

1. On zookeeper nodes:

```
service mapr-zookeeper start
```

2. On all nodes, including zookeeper nodes:

```
service mapr-warden start
```

3. Wait for about 5 minutes for the cluster to come up.

Add License to MapR Cluster

To add a license to the MapR Cluster, complete the following steps:

1. When MapR Cluster is installed and running, log in to any MapR cluster node and retrieve your unique MapR cluster ID:

```
maprcli license showid
```

2. Send the cluster ID and the number of MapR nodes, along with <customer info> to <MapR> to get your MapR License.
3. If your cluster has internet access, install your license directly from the MapR License Server. Choose Admin from the drop-down list and choose Cluster Setting. In the Admin/ Cluster Setting click License.

The screenshot shows the MapR Admin interface for Cluster Settings. The 'Licenses' tab is selected. On the left, a table displays cluster information:

Type	Maximum	Available	Used
Server Nodes	4	0	4
Client Nodes - Basic	10	-	-
Client Nodes - PACC	10	-	-

On the right, the Cluster ID is 1432193066317607737. Below this, there are three buttons: 'Import License', 'Upload License File', and 'Copy/Paste License'. A message states: 'MapR's free trial license will allow you to get a license to try one of MapR's many services for a short period.' with a 'Get a Free Trial License' button.

4. Choose Import License, enter the credentials for mapr.com Account to retrieve the license.

Import License ×

To be able to retrieve your license information, please enter your log in credentials for My Account on mapr.com.

User Name

Password

- Alternatively, if your cluster does not have internet access, you can copy your license file to any MapR cluster node, and install it via the command line interface:

```
maprcli license add -is_file true -license <license_file>
```



The MapR license file is a text file starting with the line "-----BEGIN SIGNED MESSAGE-----" and ending with the line "-- ---END MESSAGE HASH-----". You can also use "Add licenses via upload" or "Add licenses via copy/paste" in the "Manage Licenses" dialog in the MapR Control System GUI to install the license from a file.

Create Virtual IPs

In order to create virtual IP addresses for the MapR NFS server, MAC address of HANA-Storage vNIC on all the MapR nodes, IP address range for Virtual IP and Network subnet are required.

From each MapR Server execute `ifconfig eth2` command to get the MAC Address for Ethernet device carrying VLAN for HANA-Storage.

```
eth2: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    inet 192.168.110.101 netmask 255.255.255.0 broadcast 192.168.110.255
    ether 00:25:b5:7b:00:25 txqueuelen 1000 (Ethernet)
```

Below is the command to create Virtual IP on the MapR cluster:

```
maprcli virtualip add -macs <mac_addr_node1> <mac_addr_node2> <mac_addr_nodeN> -
virtualipend <<last_virtual_ip>> -netmask <<subnet_mask>> -virtualip<<first_virtual_ip>>
```

Example:

```
maprcli virtualip add -macs 00:25:b5:fb:00:29 00:25:b5:fb:00:22 00:25:b5:fb:00:24
00:25:b5:fb:00:26 -virtualipend 192.168.110.78 -netmask 255.255.255.0 -virtualip
192.168.110.71
```

Create Volumes

The MapR cluster needs an internal directory structure for its volumes. It is recommended to store the volumes in the /apps folder of the MapR cluster directory structure. The MapR file system is mounted on the MapR cluster nodes under /mapr, so the following directory structure should be created on one of the MapR cluster node.

1. Confirm MapR file system is mounted at /mapr:

```
mount | grep mapr
localhost:/mapr on /mapr type nfs (rw,nolock,addr=127.0.0.1)
```

2. Create HANA directory:

```
mkdir -p /mapr/<clustername>/apps/hana/<SID>
```

Example:

```
mkdir -p /mapr/mapr600/apps/hana/ANA
```



When creating a volume, the volume name can differ from the directory name of MapR internal file system.

3. Create volume for /hana/shared:

```
maprcli volume create -name <volumename> -path /apps/hana/<SID>/<volumename> -
replicationtype high_throughput -replication 2 -minreplication 2
```

4. Create volumes for /hana/data for every single storage partition:

```
maprcli volume create -name <volumename> -path
/apps/hana/<SID>/<datamountpoint_storage_partition> -replicationtype
high_throughput -replication 2 -minreplication 2
```

5. Create volumes for /hana/log for every single storage partition:

```
maprcli volume create -name <volumename> -path
/apps/hana/<SID>/<logmountpoint_storage_partition> -replicationtype low_latency -
replication 2 -minreplication 2
```

6. Turn off the compression for log volumes:

```
hadoop mfs -setcompression off /mapr/<clustername>/apps/hana/<SID>/<volumename>
```



It is recommended to reflect the SID in the path. For the three different volume types shared, data and log.

7. Example for creating a volume for /hana/shared for SID "ANA":

```
maprcli volume create -name shared -path /apps/hana/ANA/shared -replicationtype
high_throughput -replication 2 -minreplication 2
```

8. Example for creating eight volumes for /hana/data for SID "ANA":

```
maprcli volume create -name data001 -path /apps/hana/ANA/data001 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data002 -path /apps/hana/ANA/data002 -replicationtype
high_throughput -replication 2 -minreplication 2
```

```
maprcli volume create -name data003 -path /apps/hana/ANA/data003 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data004 -path /apps/hana/ANA/data004 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data005 -path /apps/hana/ANA/data005 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data006 -path /apps/hana/ANA/data006 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data007 -path /apps/hana/ANA/data007 -replicationtype
high_throughput -replication 2 -minreplication 2
maprcli volume create -name data008 -path /apps/hana/ANA/data008 -replicationtype
high_throughput -replication 2 -minreplication 2
```

9. Example for creating eight volumes for /hana/log for SID "ANA" :

```
maprcli volume create -name log001 -path /apps/hana/ANA/log001 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log002 -path /apps/hana/ANA/log002 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log003 -path /apps/hana/ANA/log003 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log004 -path /apps/hana/ANA/log004 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log005 -path /apps/hana/ANA/log005 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log006 -path /apps/hana/ANA/log006 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log007 -path /apps/hana/ANA/log007 -replicationtype
low_latency -replication 2 -minreplication 2
maprcli volume create -name log008 -path /apps/hana/ANA/log008 -replicationtype
low_latency -replication 2 -minreplication 2
```

10. Example to turn off compression on eight log volumes for SID "ANA" :

```
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log001
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log002
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log003
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log004
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log005
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log006
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log007
hadoop mfs -setcompression off /mapr/mapr600/apps/hana/ANA/log008
```



MapR FS is a distributed File System and there should be a quota set for each volume to restrict the volume size: Hard quota, which raises an alarm when the threshold is reached and prevents further writes. Advisory quota, which raises an alarm when the threshold is reached, but does not prevent further writes. Every volume will be set with quota as per SAP's recommended volume size. Advisory quota will be set at 85% of recommended volume size.

11. Shared filesystem should be 1 x Memory for every 4 servers. Set Quota for /hana/shared for SID ANA" :

```
maprcli volume modify -name shared -advisoryquota 5223G -quota 6144G
```

12. Data volume should be 3 x Memory. Set Quota for eight data volumes for SID "ANA" :

```
maprcli volume modify -name data001 -advisoryquota 7834G -quota 9216G
maprcli volume modify -name data002 -advisoryquota 7834G -quota 9216G
maprcli volume modify -name data003 -advisoryquota 7834G -quota 9216G
maprcli volume modify -name data004 -advisoryquota 7834G -quota 9216G
maprcli volume modify -name data005 -advisoryquota 7834G -quota 9216G
maprcli volume modify -name data006 -advisoryquota 7834G -quota 9216G
```

```
maprcli volume modify -name data007 -advisoryquota 7834G -quota 9216G  
maprcli volume modify -name data008 -advisoryquota 7834G -quota 9216G
```

13. Log volume should be 512 GB. Set Quota for eight log volumes for SID "ANA" :

```
maprcli volume modify -name log001 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log002 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log003 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log004 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log005 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log006 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log007 -advisoryquota 435G -quota 512G  
maprcli volume modify -name log008 -advisoryquota 435G -quota 512G
```

HANA System Configuration

This section provides the procedure for RedHat Enterprise Linux 7.4 Operating System and customizing for SAP HANA requirement.




For the latest information on SAP HANA installation and OS customization requirement, see the [SAP HANA Installation Guide](#).

RHEL Operating System Installation

To install the SUSE Linux Enterprise Server for SAP Applications 12 SP 2 on the local drives, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Service Profile > root > Sub-Organization > HANA > HANA-Server01.
3. Click KVM Console.
4. When the KVM Console is launched, click Boot Server.
5. Click Virtual Media > Activate Virtual Devices.
 - a. Choose the option Accept this Session for Unencrypted Virtual Media Session and then click Apply.
 - b. Click Virtual Media and Choose Map CD/DVD.
 - c. Click Browse to navigate ISO media location.
 - d. Click Map Device.
6. On the Initial screen choose Install to begin the installation process.
7. Choose Language, and click Continue.
8. This brings us to the central Installation Summary page where we need to configure various features



9. Click Date & Time; choose the appropriate timezone and click Done.
10. Click Keyboard; choose Keyboard layout and click Done.
11. Under Software Menu, click Software selection.
12. In the Base Environment choose Infrastructure Server.
13. For Add-Ons for Selected Environment choose Large Systems Performance, Network File System Client, Performance Tools, Compatibility Libraries and click Done.

SOFTWARE SELECTION RED HAT ENTERPRISE LINUX 7.4 INSTALLATION

Done us Help

Base Environment

- Minimal Install**
Basic functionality.
- Infrastructure Server**
Server for operating network infrastructure services.
- File and Print Server**
File, print, and storage server for enterprises.
- Basic Web Server**
Server for serving static and dynamic internet content.
- Virtualization Host**
Minimal virtualization host.
- Server with GUI**
Server for operating network infrastructure services, with a GUI.

Add-Ons for Selected Environment

- RDMA-based InfiniBand and iWARP fabrics.
- Java Platform**
Java support for the Red Hat Enterprise Linux Server and Desktop Platforms.
- Large Systems Performance**
Performance support tools for large systems.
- Load Balancer**
Load balancing support for network traffic.
- MariaDB Database Server**
The MariaDB SQL database server, and associated packages.
- Network File System Client**
Enables the system to attach to network storage.
- Performance Tools**
Tools for diagnosing system and application-level performance problems.
- PostgreSQL Database Server**
The PostgreSQL SQL database server, and associated packages.
- Print Server**
Allows the system to act as a print server.
- Remote Management for Linux**
Remote management interface for Red Hat Enterprise Linux, including OpenLM and SNMP.
- Virtualization Hypervisor**
Smallest possible virtualization host installation.
- Compatibility Libraries**
Compatibility libraries for applications built on previous versions of Red Hat Enterprise Linux.
- Development Tools**
A basic development environment.
- Security Tools**
Security tools for integrity and trust verification.
- Smart Card Support**
Support for using smart card authentication.

14. Under System; click Installation destination.

INSTALLATION DESTINATION
RED HAT ENTERPRISE LINUX 7.4 INSTALLATION

Done
Help!

Device Selection
 Select the device(s) you'd like to install to. They will be left untouched until you click on the main menu's "Begin Installation" button.

Local Standard Disks

278.46 GiB



Cisco UCSC-RAID12GP-4G

sda / 176.27 GiB free

Disks left unselected here will not be touched.

Specialized & Network Disks



Add a disk...

Disks left unselected here will not be touched.

Other Storage Options

Partitioning

Automatically configure partitioning.
 I will configure partitioning.

I would like to make additional space available.

Encryption


Encrypt my data. You'll set a passphrase next.

[Full disk summary and boot loader...](#)
1 disk selected; 278.46 GiB capacity; 176.27 GiB free [Refresh...](#)

15. Choose the Boot drives with RAID 1 created.

16. From the Other Storage Options choose 'I will configure partitioning' and click Done.


17. In the Manual Partitioning Screen, choose Standard Partition for New mount points will use the following partitioning scheme.




18. Click the + symbol to add a new partition.

19. Choose the mount point as **/boot**.

20. Enter the Desired capacity as 200 MB and click Add Mount Point.



21. Choose the filesystem ext3.
22. Click the + symbol to add a new partition.
23. Choose the mount point swap.
24. Enter the Desired capacity 2048 MB and click Add Mount Point.
25. Choose the filesystem swap.
26. Click the + symbol to add / (root) partition.
27. Choose the mount point as /.
28. Enter the Desired capacity 102400 MiB and click Add Mount Point.
29. Choose the filesystem ext3.




30. Click Done on the top left corner of the screen.
31. Review the partition layout and the size.
32. Click Accept Changes to proceed with the next steps.
33. Click KDUMP from the below screen.
34. Deselect Enable kdump.
35. Click Security policy, choose Apply Security policy to OFF and click Done.
36. Click Done on the top left corner of the screen.
37. Click Network & Hostname.
38. Enter the Host name and click Apply.

 Starting Red Hat Enterprise Linux 7.x, the system assigns the Ethernet device names based on the bus where they are present. The default Ethernet device naming convention like eth0, eth1 ...ethx is deprecated. This can be addressed once the OS is installed and active.

39. IP address will be assigned once the OS is installed. Click Done


40. Click Done on the top left corner of the screen.

41. Review the installation summary and click Begin Installation.



42. The next screen will show the start of the OS installation.

43. Click Root Password.



44. Enter the Root Password and Confirm.

45. Click Done on the top left corner of the window.

46. The installation will start and continue.

47. When the installation is complete click Reboot to finish the installation.

Operating System Network Configuration

To customize the server in preparation for the HANA Installation, complete the following steps:

Hostnames



The operating system must be configured in such a way that the short name of the server is displayed for the command **hostname -s**, and the fully qualified host name is displayed with the command **hostname -f**

1. Use the KVM console to log in to the installed system as the user root and the password <<var_sys_root-pw>>.
2. Set the hostname using `hostnamectl`.
3. `hostnamectl set-hostname <<hostname>>`.

```
hostnamectl set-hostname <<hostname>>
```

4. Edit the Hostname:

```
vi /etc/HOSTNAME
<<hostname>>.<<Domain Name>>
```

IP Address

Each SAP HANA Server is configured with 9 vNIC device. Table 12 lists the IP Address information required to configure the IP address on the Operating System.



The IP Address and Subnet Mask provided below are examples only, please configure the IP address for your environment.

Table 12 IP Addresses for SAP HANA Server

vNIC Name	VLAN ID	IP Address Range	Subnet Mask
HANA-AppServer	<<var_appserver_vlan_id>>	192.168.223.211 to 192.168.223.218	255.255.255.0
HANA-Backup	<<var_backup_vlan_id>>	192.168.221.211 to 192.168.221.218	255.255.255.0
HANA-Client	<<var_client_vlan_id>>	192.168.222.211 to 192.168.222.218	255.255.0.0
HANA-DataSource	<<var_datasource_vlan_id>>	192.168.224.211 to 192.168.224.218	255.255.255.0
HANA-Internal	<<var_internal_vlan_id>>	192.168.220.211 to 192.168.220.218	255.255.255.0
HANA-Replication	<<var_replication_vlan_id>>	192.168.225.211 to 192.168.225.218	255.255.255.0
HANA-Storage	<<var_storage_vlan_id>>	192.168.110.211 to 192.168.110.218	255.255.255.0
Management	<<var_mgmt_vlan_id>>	192.168.196.211 to 192.168.196.218	255.255.0.0
NFS-IPMI	<<var_inband_vlan_id>>	192.168.197.211 to 192.168.197.218	255.255.255.0

5. To configure the network interface on the OS, it is required to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS.

6. In RHEL 7, systemd and udev support a number of different naming schemes. By default, fixed names **are assigned based on firmware, topology, and location information, like 'enp72s0'**. With this naming convention, though names stay fixed even if hardware is added or removed it is often harder to read unlike traditional kernel-native ethX naming "eth0". Another way to name network interfaces, "biosdev-names", is already available with installation.
7. Configure boot parameters "net.ifnames=0 biosdevname=0" to disable both, to get the original kernel native network names.
8. Also, IPV6 support could be disabled at this time as we use IPV4 in the solution. This can be done by appending ipv6.disable=1 to GRUB_CMDLINE_LINUX as shown below:

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

9. To Run the grub2-mkconfig command to regenerate the grub.cfg file:

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

10. Finally reboot system to effect the changes.
11. To configure the network interface on the OS, it is required to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS.
12. From the OS, execute the following command to get list of Ethernet device with MAC Address

```
[root@cishanaso11 ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:05 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:04 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:06 brd ff:ff:ff:ff:ff:ff
5: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:05 brd ff:ff:ff:ff:ff:ff
6: eth4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:07 brd ff:ff:ff:ff:ff:ff
7: eth5: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:06 brd ff:ff:ff:ff:ff:ff
8: eth6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:08 brd ff:ff:ff:ff:ff:ff
9: eth7: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7b:00:07 brd ff:ff:ff:ff:ff:ff
10: eth8: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
   link/ether 00:25:b5:7a:00:09 brd ff:ff:ff:ff:ff:ff
```

13. In Cisco UCS Manager, click the Servers tab in the navigation pane.
14. Expand Servers > Service Profile > root > Sub-Organization > HANA > HANAServer01.
15. Click + to Expand.

16. Click vNICs.

17. On the right pane list of the vNICs with MAC Address are listed.

vNICs

Advanced Filter
 Export
 Print

Name	MAC Address	Desired Order	Actual Order	Fabric ID
vNIC HANA-Internal	00:25:B5:7A:00:05	1	1	A B
vNIC HANA-AppSer...	00:25:B5:7A:00:06	3	3	A B
vNIC Mgmt	00:25:B5:7A:00:07	5	5	A B
vNIC HANA-DataSo...	00:25:B5:7A:00:08	7	7	A B
vNIC NFS-IPMI	00:25:B5:7A:00:09	9	9	A B
vNIC HANA-Storage	00:25:B5:7B:00:04	2	2	B A

Delete
 Add
 Modify

18. Note the MAC Address of the HANA-Internal vNIC is “00:25:B5:7A:00:05”.

19. By comparing MAC Address on the OS and Cisco UCS, eth0 on OS will carry the VLAN for HANA-Internal.

20. Go to network configuration directory and create a configuration for eth0:

```
cd /etc/sysconfig/network-scripts/
vi ifcfg-eth0
DEVICE=eth0
TYPE=Ethernet
ONBOOT=yes
BOOTPROTO=static
IPV6INIT=no
USERCTL=no
NM_CONTROLLED=no
IPADDR=<<IP address for HANA-Internal network example:192.168.220.211>>
NETMASK=<<subnet mask for HANA-Internal network 255.255.255.0>>
```

21. Repeat the steps 11 to 19 for each vNIC interface.

22. Add default gateway:

```
vi /etc/sysconfig/network

NETWORKING=yes
HOSTNAME=<<HOSTNAME>>
GATEWAY=<<IP Address of default gateway>>
```

DNS

Domain Name Service configuration must be done based on the local requirements.

Configuration Example

Add DNS IP if it is required to access internet:

```
vi /etc/resolv.conf

DNS1=<<IP of DNS Server1>>
DNS2=<<IP of DNS Server2>>
DOMAIN= <<Domain_name>>
```

Hosts file

For scale-out system all nodes should be able to resolve internal network IP address, below is an example of 8 node host file with the entire network defined in the /etc/hosts file.

```
cat /etc/hosts
```

```
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
#
## Internal Network
#
192.168.220.211  cishanaso11.ciscolab.local cishanaso11
192.168.220.212  cishanaso12.ciscolab.local cishanaso12
192.168.220.213  cishanaso13.ciscolab.local cishanaso13
192.168.220.214  cishanaso14.ciscolab.local cishanaso14
192.168.220.215  cishanaso15.ciscolab.local cishanaso15
192.168.220.216  cishanaso16.ciscolab.local cishanaso16
192.168.220.217  cishanaso17.ciscolab.local cishanaso17
192.168.220.218  cishanaso18.ciscolab.local cishanaso18
#
## Storage Network
#
192.168.110.211  cishanaso11s.ciscolab.local cishanaso11s
192.168.110.212  cishanaso12s.ciscolab.local cishanaso12s
192.168.110.213  cishanaso13s.ciscolab.local cishanaso13s
192.168.110.214  cishanaso14s.ciscolab.local cishanaso14s
192.168.110.215  cishanaso15s.ciscolab.local cishanaso15s
192.168.110.216  cishanaso16s.ciscolab.local cishanaso16s
192.168.110.217  cishanaso17s.ciscolab.local cishanaso17s
192.168.110.218  cishanaso18s.ciscolab.local cishanaso18s
#
## Client Network
#
192.168.222.211  cishanaso11c.ciscolab.local cishanaso11c
192.168.222.212  cishanaso12c.ciscolab.local cishanaso12c
192.168.222.213  cishanaso13c.ciscolab.local cishanaso13c
192.168.222.214  cishanaso14c.ciscolab.local cishanaso14c
192.168.222.215  cishanaso15c.ciscolab.local cishanaso15c
192.168.222.216  cishanaso16c.ciscolab.local cishanaso16c
192.168.222.217  cishanaso17c.ciscolab.local cishanaso17c
192.168.222.218  cishanaso18c.ciscolab.local cishanaso18c
#
## AppServer Network
#
192.168.223.211  cishanaso11a.ciscolab.local cishanaso11a
192.168.223.212  cishanaso12a.ciscolab.local cishanaso12a
192.168.223.213  cishanaso13a.ciscolab.local cishanaso13a
192.168.223.214  cishanaso14a.ciscolab.local cishanaso14a
192.168.223.215  cishanaso15a.ciscolab.local cishanaso15a
192.168.223.216  cishanaso16a.ciscolab.local cishanaso16a
192.168.223.217  cishanaso17a.ciscolab.local cishanaso17a
```

```

192.168.223.218 cishanaso18a.ciscolab.local cishanaso18a
#
## Management Network
#
192.168.73.211 cishanaso11m.ciscolab.local cishanaso11m
192.168.73.212 cishanaso12m.ciscolab.local cishanaso12m
192.168.73.213 cishanaso13m.ciscolab.local cishanaso13m
192.168.73.214 cishanaso14m.ciscolab.local cishanaso14m
192.168.73.215 cishanaso15m.ciscolab.local cishanaso15m
192.168.73.216 cishanaso16m.ciscolab.local cishanaso16m
192.168.73.217 cishanaso17m.ciscolab.local cishanaso17m
192.168.73.218 cishanaso18m.ciscolab.local cishanaso18m
#
## Backup Network
#
192.168.221.211 cishanaso11b.ciscolab.local cishanaso11b
192.168.221.212 cishanaso12b.ciscolab.local cishanaso12b
192.168.221.213 cishanaso13b.ciscolab.local cishanaso13b
192.168.221.214 cishanaso14b.ciscolab.local cishanaso14b
192.168.221.215 cishanaso15b.ciscolab.local cishanaso15b
192.168.221.216 cishanaso16b.ciscolab.local cishanaso16b
192.168.221.217 cishanaso17b.ciscolab.local cishanaso17b
192.168.221.218 cishanaso18b.ciscolab.local cishanaso18b
#
## DataSource Network
#
192.168.224.211 cishanaso11d.ciscolab.local cishanaso11d
192.168.224.212 cishanaso12d.ciscolab.local cishanaso12d
192.168.224.213 cishanaso13d.ciscolab.local cishanaso13d
192.168.224.214 cishanaso14d.ciscolab.local cishanaso14d
192.168.224.215 cishanaso15d.ciscolab.local cishanaso15d
192.168.224.216 cishanaso16d.ciscolab.local cishanaso16d
192.168.224.217 cishanaso17d.ciscolab.local cishanaso17d
192.168.224.218 cishanaso18d.ciscolab.local cishanaso18d
#
## Replication Network
#
192.168.225.211 cishanaso11r.ciscolab.local cishanaso11r
192.168.225.212 cishanaso12r.ciscolab.local cishanaso12r
192.168.225.213 cishanaso13r.ciscolab.local cishanaso13r
192.168.225.214 cishanaso14r.ciscolab.local cishanaso14r
192.168.225.215 cishanaso15r.ciscolab.local cishanaso15r
192.168.225.216 cishanaso16r.ciscolab.local cishanaso16r
192.168.225.217 cishanaso17r.ciscolab.local cishanaso17r
192.168.225.218 cishanaso18r.ciscolab.local cishanaso18r
#
## IPMI OS Address
#
192.168.197.211 cishanaso11o-ipmi
192.168.197.212 cishanaso12o-ipmi
192.168.197.213 cishanaso13o-ipmi
192.168.197.214 cishanaso14o-ipmi
192.168.197.215 cishanaso15o-ipmi
192.168.197.216 cishanaso16o-ipmi
192.168.197.217 cishanaso17o-ipmi
192.168.197.218 cishanaso18o-ipmi
#
## IPMI External Address
#
192.168.73.134 cishanaso11-ipmi
192.168.73.135 cishanaso12-ipmi
192.168.73.136 cishanaso13-ipmi
192.168.73.137 cishanaso14-ipmi
192.168.73.138 cishanaso15-ipmi
192.168.73.139 cishanaso16-ipmi
192.168.73.140 cishanaso17-ipmi
192.168.73.141 cishanaso18-ipmi
#
## NFS-IPMI Inband Address
#
192.168.197.164 cishanaso11i-ipmi

```

```

192.168.197.165 cishanaso12i-ipmi
192.168.197.166 cishanaso13i-ipmi
192.168.197.167 cishanaso14i-ipmi
192.168.197.168 cishanaso15i-ipmi
192.168.197.169 cishanaso16i-ipmi
192.168.197.170 cishanaso17i-ipmi
192.168.197.171 cishanaso18i-ipmi
#
##MapR Storage
#
192.168.110.101 mapr11s.ciscolab.local mapr11s
192.168.110.102 mapr12s.ciscolab.local mapr12s
192.168.110.103 mapr13s.ciscolab.local mapr13s
192.168.110.104 mapr14s.ciscolab.local mapr14s
#
## MapR Storage Virtual IP
#
192.168.110.71 maprvip31
192.168.110.72 maprvip32
192.168.110.73 maprvip33
192.168.110.74 maprvip34
192.168.110.75 maprvip35
192.168.110.76 maprvip36
192.168.110.77 maprvip37
192.168.110.78 maprvip38

```



To get the IP address assigned to the Cisco UCS service profile for Inband Management, go to Cisco UCS Manager, click the LAN tab in the navigation pane, choose Pools > root > Sub-Organization > HANA > IP Pools > IP Pools HANA-Servers. In the right pane click the IP Address tab that shows the IP address and assigned Service Profiles.

RedHat System Update and OS customization for SAP HANA

To update and customize SAP HANA, complete the following steps:

1. Updating the RedHat system.



In order to patch the system, the repository must be updated. Note that the installed system does not include any update information. In order to patch the RedHat System, it must be registered and attached to a valid Subscription. The following line will register the installation and update the repository information.

```

subscription-manager register --auto-attach
Username: <<username>>
Password: <<password>>

```

2. Update only the OS kernel and firmware packages to the latest release that appeared in RHEL 7.4. Set the release version to 7.4:

```
subscription-manager release --set=7.4
```

3. Add the repos required for SAP HANA:

```

subscription-manager repos --disable "*"
subscription-manager repos --enable rhel-7-server-rpms --enable rhel-sap-hana-for-rhel-7-server-rpms

```

4. Apply the latest updates for RHEL 7.4 Typically, the kernel is updated as well:

```
yum -y update
```

5. Install dependencies in accordance with the SAP HANA Server Installation and Update Guide. The `numactl` package if the benchmark HWCCT is to be used:

```
yum -y install gtk2 libicu xulrunner sudo tcsh libssh2 expect cairo graphviz iptraf krb5-workstation libpng12
krb5-libs nfs-utils lm_sensors rsyslog compat-sap-c++-* openssl098e openssl PackageKit-gtk-module
libcanberra-gtk2 libtool-ltdl xorg-x11-xauth compat-libstdc++-33 numactl libuuid uuid e2fsprogs icedtea-web
xfspgms net-tools bind-utils glibc-devel libgomp
```

6. Disable SELinux:

```
sed -i 's/^SELINUX=enforcing/SELINUX=disabled/g' /etc/sysconfig/selinux
sed -i 's/^SELINUX=permissive/SELINUX=disabled/g' /etc/sysconfig/selinux
sed -i 's/^SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
sed -i 's/^SELINUX=permissive/SELINUX=disabled/g' /etc/selinux/config
```

7. Sysctl.conf: The following parameters must be set in `/etc/sysctl.conf`:

```
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.conf.all.rp_filter = 0
net.ipv4.ip_local_port_range = 40000 61000
net.ipv4.neigh.default.gc_thresh1 = 256
net.ipv4.neigh.default.gc_thresh2 = 1024
net.ipv4.neigh.default.gc_thresh3 = 4096
net.ipv6.neigh.default.gc_thresh1 = 256
net.ipv6.neigh.default.gc_thresh2 = 1024
net.ipv6.neigh.default.gc_thresh3 = 4096
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.rmem_default = 262144
net.core.wmem_default = 262144
net.core.optmem_max = 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_rmem = 65536 262144 16777216
net.ipv4.tcp_wmem = 65536 262144 16777216
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
sunrpc.tcp_max_slot_table_entries = 128
```

8. Add the following line into `/etc/modprobe.d/sunrpc-local.conf`. Create the file, if it does not exist:

```
sunrpc.tcp_max_slot_table_entries = 128
```

9. For compatibility reasons, four symbolic links are required:

```
# ln -s /usr/lib64/libssl.so.0.9.8e /usr/lib64/libssl.so.0.9.8
# ln -s /usr/lib64/libssl.so.1.0.1e /usr/lib64/libssl.so.1.0.1
# ln -s /usr/lib64/libcrypto.so.0.9.8e /usr/lib64/libcrypto.so.0.9.8
# ln -s /usr/lib64/libcrypto.so.1.0.1e /usr/lib64/libcrypto.so.1.0.1
```

10. The Linux kernel shipped with RHEL 7 includes a `cpuidle` driver for recent Intel CPUs: `intel_idle`. This driver leads to a different behavior in C-states switching. The normal operating state is C0, when the processor is put to a higher C state, which saves power. But for low latency applications, the additional

time needed to stop and start the execution of the code will cause performance degradations. Modify the file `/etc/default/grub` and append the following parameter to the line starting with `GRUB_CMDLINE_LINUX`:

```
transparent_hugepage=never intel_idle.max_cstate=1 processor.max_cstate=1
```

11. To implement these changes, rebuild the GRUB2 configuration

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

12. Turn off auto-numa balancing: SAP HANA is a NUMA (non-uniform memory access) aware database. Thus it does not rely on the Linux kernel's features to optimize NUMA usage automatically. Depending on the workload, it can be beneficial to turn off automatic NUMA balancing. For this purpose, add `kernel.numa_balancing = 0` to `/etc/sysctl.d/sap_hana.conf` (please create this file if it does not already exist) and reconfigure the kernel by running:

```
#echo "kernel.numa_balancing = 0" >> /etc/sysctl.d/sap_hana.conf
#sysctl -p /etc/sysctl.d/sap_hana.conf
```

13. The "numad" daemon must be disable:

```
#systemctl stop numad
#systemctl disable numad
```

14. Configure tuned to use profile "sap-hana". The tuned profile "sap-hana", which is provided by Red Hat as part of RHEL 7 for SAP HANA, contains many of the configures some additional settings. Therefore the "sap-hana" tuned profile must be activated on all systems running SAP HANA:

```
# yum install tuned-profiles-sap-hana
# systemctl start tuned
# systemctl enable tuned
# tuned-adm profile sap-hana
```

15. Disable ABRT, Crash Dump:

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

16. Disable core file creation. To disable core dumps for all users, open `/etc/security/limits.conf`, and add the line:

```
* soft core 0
* hard core 0
```

17. Enable group "sapsys" to create an unlimited number of processes:

```
echo "@sapsys soft nproc unlimited" > /etc/security/limits.d/99-sapsys.conf
```

18. Disable Firewall:

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

19. Reboot the OS by issuing `reboot` command.
20. Optional: old kernels can be removed after OS update:

```
package-cleanup --oldkernels --count=1 -y
```

Install Cisco eNIC Driver

To download the Cisco UCS Drivers ISO bundle, which contains most of the Cisco UCS Virtual Interface Card drivers, complete the following steps:

1. In a web browser, navigate to <https://www.cisco.com>.
2. Under Support, click All Downloads.
3. In the product selector, click Products, then click Server - Unified Computing.
4. If prompted, enter your Cisco.com username and password to log in.



You must be signed in to download Cisco Unified Computing System (UCS) drivers.

5. Cisco UCS drivers are available for both Cisco UCS B-Series Blade Server Software and Cisco UCS C-Series Rack-Mount UCS-Managed Server Software.
6. Click UCS B-Series Blade Server Software.
7. Click Cisco Unified Computing System (UCS) Drivers.



The latest release version is selected by default. This document is built on Version 3.2(2a).

8. Click 3.2(2a) Version.
9. Download ISO image of Cisco Unified Computing System (UCS) Drivers.
10. Choose ISO image of UCS-related linux drivers only and click Download and follow the prompts to complete your driver download.
11. After the download is complete browse to `ucs-cxxx-drivers-linux.3.1.2b\Network\Cisco\VIC\RHEL\RHEL7.4` and copy `kmod-enic-2.3.0.44-rhel7u4.e17.x86_64.rpm` to each server.
12. ssh to the Server on the Management IP as root.
13. Update the enic driver with the following command:

```
rpm -Uvh /tmp/kmod-enic-2.3.0.44-rhel7u4.e17.x86_64.rpm
```

14. Update the enic driver on all the HANA Servers.

Network Time

The configuration of NTP is important and must be performed on all systems. To configure network time, complete the following steps:

1. Install NTP-server with utilities:

```
yum -y install ntp ntpdate
```

2. Configure NTP by adding at least one NTP server to the NTP config file /etc/ntp.conf:

```
vi /etc/ntp.conf
server <NTP-SERVER1 IP>
server <NTP-SERVER2 IP>
```

3. Stop the NTP services and update the NTP Servers:

```
systemctl stop ntpd
ntpdate ntp.example.com
```

4. Start NTP service and configure it to be started automatically:

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

SSH Keys

The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system for user 'root' and user <SID>adm.

1. Generate the rsa public key by executing the command:

```
ssh-keygen -b 2048
```

2. The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system for user 'root' and user.
3. Exchange the rsa public key by executing the following command from the first server to rest of the servers in the scale-out system:

```
ssh-copy-id -i /root/.ssh/id_rsa.pub cishanaso12
```

4. Repeat the 1- 3 for all the servers in the SAP HANA system.

Mount Options

1. Mount options vary from the default Linux setting for using NFS for SAP HANA data and log volumes. The following is an example of /etc/fstab entry for an eight node SAP HANA Scale-Out system with SID ANA:

```
#HANA Shared Volume
maprvip31:/mapr/mapr600/apps/hana/ANA/shared /hana/shared nfs nolock,hard,timeo=600 0 0
#HANA Data Volume
maprvip31:/mapr/mapr600/apps/hana/ANA/data001 /hana/data/ANA/mnt00001 nfs nolock,hard,timeo=600 0 0
```

```

maprvip32:/mapr/mapr600/apps/hana/ANA/data002 /hana/data/ANA/mnt00002 nfs nolock,hard,timeo=600 0 0
maprvip33:/mapr/mapr600/apps/hana/ANA/data003 /hana/data/ANA/mnt00003 nfs nolock,hard,timeo=600 0 0
maprvip34:/mapr/mapr600/apps/hana/ANA/data004 /hana/data/ANA/mnt00004 nfs nolock,hard,timeo=600 0 0
maprvip35:/mapr/mapr600/apps/hana/ANA/data005 /hana/data/ANA/mnt00005 nfs nolock,hard,timeo=600 0 0
maprvip36:/mapr/mapr600/apps/hana/ANA/data006 /hana/data/ANA/mnt00006 nfs nolock,hard,timeo=600 0 0
maprvip37:/mapr/mapr600/apps/hana/ANA/data007 /hana/data/ANA/mnt00007 nfs nolock,hard,timeo=600 0 0
maprvip38:/mapr/mapr600/apps/hana/ANA/data008 /hana/data/ANA/mnt00008 nfs nolock,hard,timeo=600 0 0
#HANA Log Volume
maprvip38:/mapr/mapr600/apps/hana/ANA/log001 /hana/log/ANA/mnt00001 nfs nolock,hard,timeo=600 0 0
maprvip37:/mapr/mapr600/apps/hana/ANA/log002 /hana/log/ANA/mnt00002 nfs nolock,hard,timeo=600 0 0
maprvip36:/mapr/mapr600/apps/hana/ANA/log003 /hana/log/ANA/mnt00003 nfs nolock,hard,timeo=600 0 0
maprvip35:/mapr/mapr600/apps/hana/ANA/log004 /hana/log/ANA/mnt00004 nfs nolock,hard,timeo=600 0 0
maprvip34:/mapr/mapr600/apps/hana/ANA/log005 /hana/log/ANA/mnt00005 nfs nolock,hard,timeo=600 0 0
maprvip33:/mapr/mapr600/apps/hana/ANA/log006 /hana/log/ANA/mnt00006 nfs nolock,hard,timeo=600 0 0
maprvip32:/mapr/mapr600/apps/hana/ANA/log007 /hana/log/ANA/mnt00007 nfs nolock,hard,timeo=600 0 0
maprvip31:/mapr/mapr600/apps/hana/ANA/log008 /hana/log/ANA/mnt00008 nfs nolock,hard,timeo=600 0 0

```

2. Create the required directory to mount /hana/shared /hana/data and /hana/log volumes. Mount all the volumes from /etc/fstab using “mount -a”. Check the status of all mounted volumes using “df -h” command:

```

maprvip31:/mapr/mapr600/apps/hana/ANA/shared 124T 144G 124T 1% /hana/shared
maprvip31:/mapr/mapr600/apps/hana/ANA/data001 124T 144G 124T 1% /hana/data/ANA/mnt00001
maprvip32:/mapr/mapr600/apps/hana/ANA/data002 124T 144G 124T 1% /hana/data/ANA/mnt00002
maprvip33:/mapr/mapr600/apps/hana/ANA/data003 124T 144G 124T 1% /hana/data/ANA/mnt00003
maprvip34:/mapr/mapr600/apps/hana/ANA/data004 124T 144G 124T 1% /hana/data/ANA/mnt00004
maprvip35:/mapr/mapr600/apps/hana/ANA/data005 124T 144G 124T 1% /hana/data/ANA/mnt00005
maprvip36:/mapr/mapr600/apps/hana/ANA/data006 124T 144G 124T 1% /hana/data/ANA/mnt00006
maprvip37:/mapr/mapr600/apps/hana/ANA/data007 124T 144G 124T 1% /hana/data/ANA/mnt00007
maprvip38:/mapr/mapr600/apps/hana/ANA/data008 124T 144G 124T 1% /hana/data/ANA/mnt00008
maprvip38:/mapr/mapr600/apps/hana/ANA/log001 124T 144G 124T 1% /hana/log/ANA/mnt00001
maprvip37:/mapr/mapr600/apps/hana/ANA/log002 124T 144G 124T 1% /hana/log/ANA/mnt00002
maprvip36:/mapr/mapr600/apps/hana/ANA/log003 124T 144G 124T 1% /hana/log/ANA/mnt00003
maprvip35:/mapr/mapr600/apps/hana/ANA/log004 124T 144G 124T 1% /hana/log/ANA/mnt00004
maprvip34:/mapr/mapr600/apps/hana/ANA/log005 124T 144G 124T 1% /hana/log/ANA/mnt00005
maprvip33:/mapr/mapr600/apps/hana/ANA/log006 124T 144G 124T 1% /hana/log/ANA/mnt00006
maprvip32:/mapr/mapr600/apps/hana/ANA/log007 124T 144G 124T 1% /hana/log/ANA/mnt00007
maprvip31:/mapr/mapr600/apps/hana/ANA/log008 124T 144G 124T 1% /hana/log/ANA/mnt00008

```

3. Change the directory permissions BEFORE installing HANA. Use the chown command after the file systems are mounted on each HANA node:

```

chmod -R 777 /hana/data
chmod -R 777 /hana/log

```

SAP HANA Installation

Please use the official SAP documentation, which describes the installation process with and without the SAP unified installer.



Read the SAP Notes before you start the installation (see [Important SAP Notes](#)). These SAP Notes contain the latest information about the installation, as well as corrections to the installation documentation.

[SAP HANA Server Installation Guide](#)

All other SAP installation and administration documentation is available here:

<http://service.sap.com/instguides>

Important SAP Notes

Read the following SAP Notes before you start the installation. These SAP Notes contain the latest information about the installation, as well as corrections to the installation documentation.

The latest SAP Notes can be found here: <https://service.sap.com/notes>.

SAP HANA IMDB Related Notes

[SAP Note 1514967](#) - SAP HANA: Central Note

[SAP Note 1523337](#) - SAP HANA Database: Central Note

[SAP Note 2000003](#) - FAQ: SAP HANA

[SAP Note 1730999](#) - Configuration changes in SAP HANA appliance

[SAP Note 1514966](#) - SAP HANA 1.0: Sizing SAP In-Memory Database

[SAP Note 1780950](#) - Connection problems due to host name resolution

[SAP Note 1743225](#) - SAP HANA: Potential failure of connections with scale out nodes

[SAP Note 1755396](#) - Released DT solutions for SAP HANA with disk replication

[SAP Note 1890444](#) - HANA system slow due to CPU power save mode

[SAP Note 1681092](#) - Support for multiple SAP HANA databases on a single SAP HANA appliance

[SAP Note 1514966](#) - SAP HANA: Sizing SAP HANA Database

[SAP Note 1637145](#) - SAP BW on HANA: Sizing SAP HANA Database

[SAP Note 1793345](#) - Sizing for Suite on HANA

Linux Related Notes

- [SAP Note 2235581](#) - SAP HANA: Supported Operating Systems
- [SAP Note 2009879](#) - SAP HANA Guidelines for RedHat Enterprise Linux (RHEL)
- [SAP Note 2292690](#) - SAP HANA DB: Recommended OS settings for RHEL 7
- [SAP Note 2228351](#) - SAP HANA Database SPS 11 revision 110 (or higher) on RHEL 6 or SLES 11
- [SAP Note 1944799](#) - SAP HANA Guidelines for SLES Operating System
- [SAP Note 2205917](#) - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
- [SAP Note 1731000](#) - Non-recommended configuration changes
- [SAP Note 2382421](#) - Optimizing the Network Configuration on HANA- and OS-Level
- [SAP Note 1557506](#) - Linux paging improvements
- [SAP Note 1740136](#) - SAP HANA: wrong mount option may lead to corrupt persistency
- [SAP Note 1829651](#) - Time zone settings in SAP HANA scale out landscapes

SAP Application Related Notes

- [SAP Note 1658845](#) - SAP HANA DB hardware check
- [SAP Note 1637145](#) - SAP BW on SAP HANA: Sizing SAP In-Memory Database
- [SAP Note 1661202](#) - Support for multiple applications on SAP HANA
- [SAP Note 1681092](#) - Support for multiple SAP HANA databases one HANA aka Multi SID
- [SAP Note 1577128](#) - Supported clients for SAP HANA 1.0
- [SAP Note 1808450](#) - Homogenous system landscape for on BW-HANA
- [SAP Note 1976729](#) - Application Component Hierarchy for SAP HANA
- [SAP Note 1927949](#) - Standard Behavior for SAP Logon Tickets
- [SAP Note 1577128](#) - Supported clients for SAP HANA
- [SAP Note 2186744](#) - FAQ: SAP HANA Parameters
- [SAP Note 2267798](#) - Configuration of the SAP HANA Database during Installation Using hdbparam
- [SAP Note 2156526](#) - Parameter constraint validation on section indices does not work correctly with hdbparam
- [SAP Note 2399079](#) - Elimination of hdbparam in HANA 2

Third Party Software

[SAP Note 1730928](#) - Using external software in a SAP HANA appliance

[SAP Note 1730929](#) - Using external tools in an SAP HANA appliance

[SAP Note 1730930](#) - Using antivirus software in an SAP HANA appliance

[SAP Note 1730932](#) - Using backup tools with Backint for SAP HANA

SAP HANA Virtualization

[SAP Note 1788665](#) - SAP HANA running on VMware vSphere VMs

High Availability (HA) Configuration for Scale-Out

For HANA Scale-Out, the ha_provider python class supports the STONITH functionality.

STONITH = Shoot the Other Node In The Head. With this python class, we are able to reboot the failing node to prevent a split brain and thus an inconsistency of the database. Since we use NFSv3, we must implement the STONITH functionality to prevent the database from a corruption because of multiple access to mounted file systems. If a HANA node is failed over to another node, the failed node will be rebooted from the master name server. This eliminates the risk of multiple accesses to the same file systems.

High Availability Configuration

1. The used version of the ucs_ha_class.py must be at least 1.1

```

vi ucs_ha_class.py
"""
Function Class to call the reset program to kill the failed host and remove NFS locks for the SAP
HANA HA
Class Name ucs_ha_class
Class Path /usr/sap/<SID>/HDB<ID>/exe/python_support/hdb_ha
Provider Cisco Systems Inc.
Version 1.1 (apiVersion=2 and hdb_ha.client import sudowers)
"""
from hdb_ha.client import StorageConnectorClient
import os
class ucs_ha_class(StorageConnectorClient):
    apiVersion = 2
    def __init__(self, *args, **kwargs):
        super(ucs_ha_class, self).__init__(*args, **kwargs)
    def stonith(self, hostname):
        os.system ("/bin/logger STONITH HANA Node:" + hostname)
        os.system ("/hana/shared/HA/ucs_ipmi_reset.sh " + hostname)
        return 0
    def about(self):
        ver={"provider_company":"Cisco",
            "provider_name" : "ucs_ha_class",
            "provider_version":"1.0",
            "api_version" :2}
        self.tracer.debug('about: %s'+str(ver))
        print '>> ha about',ver
        return ver
    @staticmethod
    def sudoers():
        return """"ALL=NOPASSWD: /bin/mount, /bin/umount, /bin/logger"""""
    def attach(self, storages):
        return 0

```

```
def detach(self, storages):
return 0
def info(self, paths):
pass
```

2. Prepare the script to match the Cisco UCS Manager configured ipmi username and password. Default is ipmi-user sapadm and ipmi-user-password cisco:

```
vi ucs_ipmi_reset.sh
#!/bin/bash
# Cisco Systems Inc.
# SAP HANA High Availability
# Version NFS_UCS v01
# changelog: 02/18/2016
if [ -z $1 ]
then
echo "please add the hostname to reset to the command line"
exit 1
fi
# Trim the domain name off of the hostname
host=`echo "$1" | awk -F'.' '{print $1}'`
PASSWD=cisco
USER=sapadm
echo $host-ipmi
system_down='Chassis Power is off'
system_up='Chassis Power is on'
power_down='power off'
power_up='power on'
power_status='power status'
#
# Shut down the server via ipmitool power off
#
/bin/logger `whoami` " Resetting the HANA Node $host because of an Nameserver reset command"
#Power Off
rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_down`
sleep 20
#Status
rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_status`
#Chassis Power is still on
if [ "$rc3" = "$system_down" ]
then
/bin/logger `whoami` " HANA Node $host switched from ON to OFF "
else
#Power Off again
/bin/logger `whoami` " HANA Node $host still online second try to shutdown... "
rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_down`
sleep 20
#Status
rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_status`
#Chassis Power is still on
if [ "$rc3" = "$system_down" ]
then
/bin/logger `whoami` " HANA Node $host switched from ON to OFF 2nd try"
else
/bin/logger `whoami` " Resetting the HANA Node $host failed "
exit 1
fi
fi
#Chassis Power is down and the server can be swiched back on
#
#The NFS locks are released
#We will start the server now to bring it back as standby node
#Chassis Power is off
power="off"
/bin/logger `whoami` " HANA Node $host will stay offline for 10 seconds..... "
sleep 10
/bin/logger `whoami` " Switching HANA Node $host back ON "
rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_up`
sleep 20
```

```

#Status
rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_status`
#Chassis Power is off
if [ "$rc3" = "$system_up" ]
then
/bin/logger `whoami` " HANA Node $host reset done, system is booting"
power="on"
exit 0
else
/bin/logger `whoami` " Switching HANA Node $host back ON failed first time..."
/bin/logger `whoami` " Switching HANA Node $host back ON second time..."
rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_up`
sleep 20
rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD $power_status`
if [ "$rc3" = "$system_up" ]
then
/bin/logger `whoami` " HANA Node $host reset done, system is booting"
power="on"
exit 0
else
/bin/logger `whoami` " Resetting the HANA Node $host failed "
exit 1
fi
fi
#
# Server is power on and should boot - our work is done
#

```

3. Copy the HA scripts to the shared HA directory under /hana/shared/<SID>/HA (HANA nameserver is responsible to reset the failed node):

```

ssh cishana01
mkdir /hana/shared/HA
chown anaadm:sapsys /hana/shared/HA
scp ucs_ipmi_reset.sh /hana/shared/HA/
scp ucs_ha_class.py /hana/shared/HA/
chown anaadm:sapsys /hana/shared/HA/*

```

Enable the SAP HANA Storage Connector API

The SAP Storage Connector API provides a way to call a user procedure whenever the SAP HANA Nameserver triggers a node failover. The API requires the files mentioned above.

The procedure is executed on the master nameserver.

1. To activate the procedure in case of a node failover, the global.ini file in:

```

<HANA installdirectory>/<SID>/global/hdb/custom/config/
must be edited and the following entry must be added:
[Storage]
ha_provider = ucs_ha_class
ha_provider_path = /hana/shared/HA
cd /hana/shared/<SID>/global/hdb/custom/config
vi global.ini
[communication]
internal_network = 192.168.220/24
listeninterface = .internal
[internal_hostname_resolution]
192.168.220.211 = cishanaso11
192.168.220.212 = cishanaso12
192.168.220.213 = cishanaso13
192.168.220.218 = cishanaso18
192.168.220.214 = cishanaso14
192.168.220.216 = cishanaso16
192.168.220.215 = cishanaso15

```

```
192.168.220.217 = cishanaso17
[persistence]
basepath_datavolumes = /hana/data/ANA
basepath_logvolumes = /hana/log/ANA
[storage]
ha_provider = ucs_ha_class
ha_provider_path = /hana/shared/HA
[trace]
ha_ucs_ha_class = info
```

2. Modify the /etc/sudoers file and append the below line on all the nodes. By adding the line <sid>adm account can execute commands mentioned without password.

```
cishana01:/ # vi /etc/sudoers
<sid>adm ALL=NOPASSWD: /bin/mount, /bin/umount, /bin/logger, /sbin/multipath,
/sbin/multipathd, /usr/bin/sg_persist, /etc/init.d/multipathd, /bin/kill,
/usr/bin/lsof, /sbin/vgchange, /sbin/vgscan
```

3. To activate the change, restart the SAP HANA DB.

Test the IPMI Connectivity

Test the ipmi connectivity on ALL nodes:

```
cishana01:~ # ipmitool -I lanplus -H cishana01-ipmi -U sapadm -P cisco power
status
Chassis Power is on
```



Make sure that all the nodes are responding to the ipmitool command and the IP address for the ipmi network match in the /etc/hosts file of all the servers.

SAP HANA Parameters for Scale-Out

1. The default communication path among the SAP HANA nodes uses the access network. To optimize the internode traffic, the MTU value needs to be 9000, which is already set in the storage VLAN. As part of setting up a distributed system, you need to configure the network parameters. Make sure that you do this before you add hosts because one server needs to be available so that you can connect to the SAP HANA Studio. For more information, refer to the SAP HANA Administration Guide.

```
hdbnsutil -reconfig --hostnameResolution=internal
```

2. After the change, all internal SAP HANA traffic goes over the server. Make sure to adjust internal_hostname_resolution in the /hana/shared/<SID>/global/hdb/custom/config/global.ini

```
[communication]
internal_network = 192.168.220/24
listeninterface = .internal
[internal_hostname_resolution]
192.168.220.211 = cishanaso11
192.168.220.212 = cishanaso12
192.168.220.213 = cishanaso13
192.168.220.218 = cishanaso18
192.168.220.214 = cishanaso14
192.168.220.216 = cishanaso16
192.168.220.215 = cishanaso15
192.168.220.217 = cishanaso17
[persistence]
basepath_datavolumes = /hana/data/ANA
```

```

basepath_logvolumes = /hana/log/ANA
[storage]
ha_provider = ucs_ha_class
ha_provider_path = /hana/shared/HA
[trace]
ha_ucs_ha_class = info

```

SAP HANA Hardware Configuration Check Tool Parameter

1. Use the following parameter setting in the HWCCT in order to meeting the SAP HANA KPI for the Filesystem Test:

```

{
  "id":2,
  "package": "FilesystemTest",
  "test_timeout": 0,
  "config": {"mount":{
    "cishanaso11":["/hana/log/ANA/mnt00001/"],
    "cishanaso12":["/hana/log/ANA/mnt00002/"],
    "cishanaso13":["/hana/log/ANA/mnt00003/"],
    "cishanaso14":["/hana/log/ANA/mnt00004/"],
    "cishanaso15":["/hana/log/ANA/mnt00005/"],
    "cishanaso16":["/hana/log/ANA/mnt00006/"],
    "cishanaso17":["/hana/log/ANA/mnt00007/"],
    "cishanaso18":["/hana/log/ANA/mnt00008/"]
  }},
  "parameter":{"max_parallel_io_requests":"128",
    "async_read_submit":"on",
    "async_write_submit_active": "on",
    "async_write_submit_blocks":"all"},
  "duration":"long"
},
"class": "LogVolumeIO"
},
{
  "id":3,
  "package": "FilesystemTest",
  "test_timeout": 0,
  "config": {"mount":{
    "cishanaso11":["/hana/data/ANA/mnt00001/"],
    "cishanaso12":["/hana/data/ANA/mnt00002/"],
    "cishanaso13":["/hana/data/ANA/mnt00003/"],
    "cishanaso14":["/hana/data/ANA/mnt00004/"],
    "cishanaso15":["/hana/data/ANA/mnt00005/"],
    "cishanaso16":["/hana/data/ANA/mnt00006/"],
    "cishanaso17":["/hana/data/ANA/mnt00007/"],
    "cishanaso18":["/hana/data/ANA/mnt00008/"]
  }},
  "parameter":{"max_parallel_io_requests":"128",
    "async_read_submit":"on",
    "async_write_submit_active": "on",
    "async_write_submit_blocks":"all"},
  "duration":"long"
},
"class": "DataVolumeIO"
}
}

```

hdbparam is used to configure the HANA database with parameter vector given for optimal storage performance. hdbparam used for HANA 1.0.

2. Run these commands for all four parameters:

```

"max_parallel_io_requests":"128",
"async_read_submit":"on",

```

```
"async_write_submit_active": "on",
"async_write_submit_blocks": "all"},
```

3. With User <sid>adm:

```
hdbparam --paramset fileio [DATA].max_parallel_io_requests=128
hdbparam --paramset fileio [DATA].async_read_submit=on
hdbparam --paramset fileio [DATA].async_write_submit_active=on
hdbparam --paramset fileio [DATA].async_write_submit_blocks=all
hdbparam --paramset fileio [LOG].max_parallel_io_requests=128
hdbparam --paramset fileio [LOG].async_read_submit=on
hdbparam --paramset fileio [LOG].async_write_submit_active=on
hdbparam --paramset fileio [LOG].async_write_submit_blocks=all
```

For SAP HANA 2.0 global.ini is used to set the parameter vector for optimal storage performance.

4. Add the following parameters in the /hana/shared/<SID>/global/hdb/custom/config/global.ini:

```
[fileio]
max_parallel_io_requests = 128
async_read_submit = on
async_write_submit_active = on
async_write_submit_blocks = all
```

5. Restart the HANA Database for configuration to take effect.

6. Stop the HANA Database with below command

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StopSystem HDB
```

7. Start the HANA Database with the following command:

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StartSystem HDB
```

8. Check the HANA Database Status with the following command:

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetSystemInstanceList
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
cishanaso16, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
cishanaso15, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
cishanaso18, 0, 50013, 50014, 0.3, HDB|HDB_STANDBY, GREEN
cishanaso11, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
cishanaso13, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
cishanaso12, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
cishanaso17, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
```

About the Authors

Shailendra Mruthunjaya, Cisco Systems, Inc.

Shailendra is a Technical Marketing Engineer with Cisco UCS Solutions and Performance Group. Shailendra has over six years of experience with SAP HANA on Cisco UCS platform. Shailendra has designed several SAP landscapes in public and private cloud environment. Currently, his focus is on developing and validating infrastructure best practices for SAP applications on Cisco UCS Servers, Cisco Nexus products and Storage technologies.

Matthias Schlarb, Cisco Systems, Inc.

Matthias is part of the Cisco SAP Competence Center in Walldorf where he conducts workshops with customers and partners. As a Technical Marketing Engineer, he develops best practices for SAP Cloud Infrastructure on Cisco Unified Computing System.

Acknowledgements

For their support and contribution to the design, validation, and creation of this CVD, we would like to thank:

- Ralf Klahr, Cisco Systems, Inc.
- Ulrich Kleidon, Cisco Systems, Inc.
- Erik Lillestolen, Cisco Systems, Inc.
- Arun Kumar Thiruttani Kalathi, Cisco Systems, Inc.
- Andy Lerner, MapR Technologies, Inc.
- James Sun, MapR Technologies, Inc.