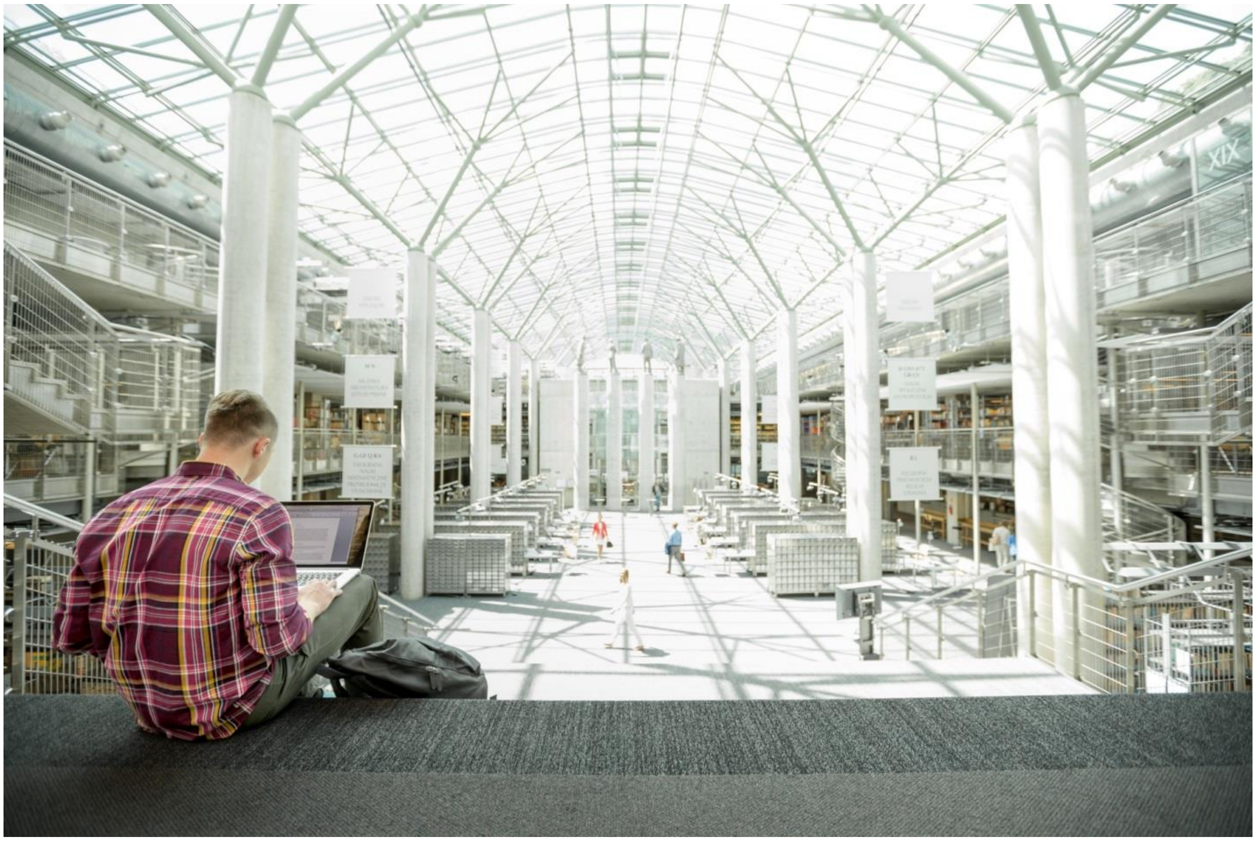

FlexPod Datacenter for SAP Solution with Cisco UCS Manager 4.0 and NetApp AFF A–Series Design Guide

Last Updated: October 4, 2019

About the Cisco Validated Design Program

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

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

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

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

Table of Contents

Executive Summary

Cisco Validated Designs (CVDs) consist of 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 our customers.

This document describes the Cisco and NetApp® FlexPod® solution, which is a validated approach for deploying SAP HANA® Tailored Data Center Integration (TDI) environments. This validated design provides guidelines and a framework for implementing SAP HANA with best practices from Cisco and NetApp.

FlexPod is a leading integrated infrastructure that supports a broad range of enterprise workloads and use cases. This solution allows you to quickly and reliably deploy SAP HANA with a model of tailored datacenter integration mode.

The recommended solution architecture is built on the Cisco Unified Computing System (Cisco UCS) using a unified software release to support Cisco UCS hardware platforms that include the following components:

- Cisco UCS B-Series blade servers and Cisco UCS C-Series rack servers configurable with Intel Optane Data Center Persistent Memory Module (DCPMM) option
- Cisco UCS 6300 series Fabric Interconnects
- Cisco Nexus 9000 Series switches
- NetApp All Flash series storage arrays

In addition, this guide provides validations for both Red Hat Enterprise Linux and SUSE Linux Enterprise Server for SAP HANA.

Solution Overview

Introduction

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

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

Audience

The audience for this document includes sales engineers, field consultants, professional services specialists, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

What's New in this Release?

This primary design for FlexPod Datacenter for SAP Solution has been updated to include the latest Cisco and NetApp hardware and software, include the following:

- Support for the Cisco UCS 4.0(4) unified software release, Cisco UCS B200-M5 servers, Cisco UCS B480-M5 servers with Cascade Lake CPUs, and Cisco 1400 Series Virtual Interface Cards (VICs)
- Validation with cloud-scale FX Series Nexus switches
- Support for the NetApp AFF A320 storage controller
- Support for the latest release of NetApp ONTAP® 9.6 storage software
- Support for NFS v4.1
- Support for NetApp SnapCenter® 4.2
- NFS and iSCSI storage design
- 100GbE connectivity

Technology Overview

FlexPod System Overview

FlexPod is a best practice datacenter architecture that is based on the following components:

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

Figure 1 FlexPod Component Families

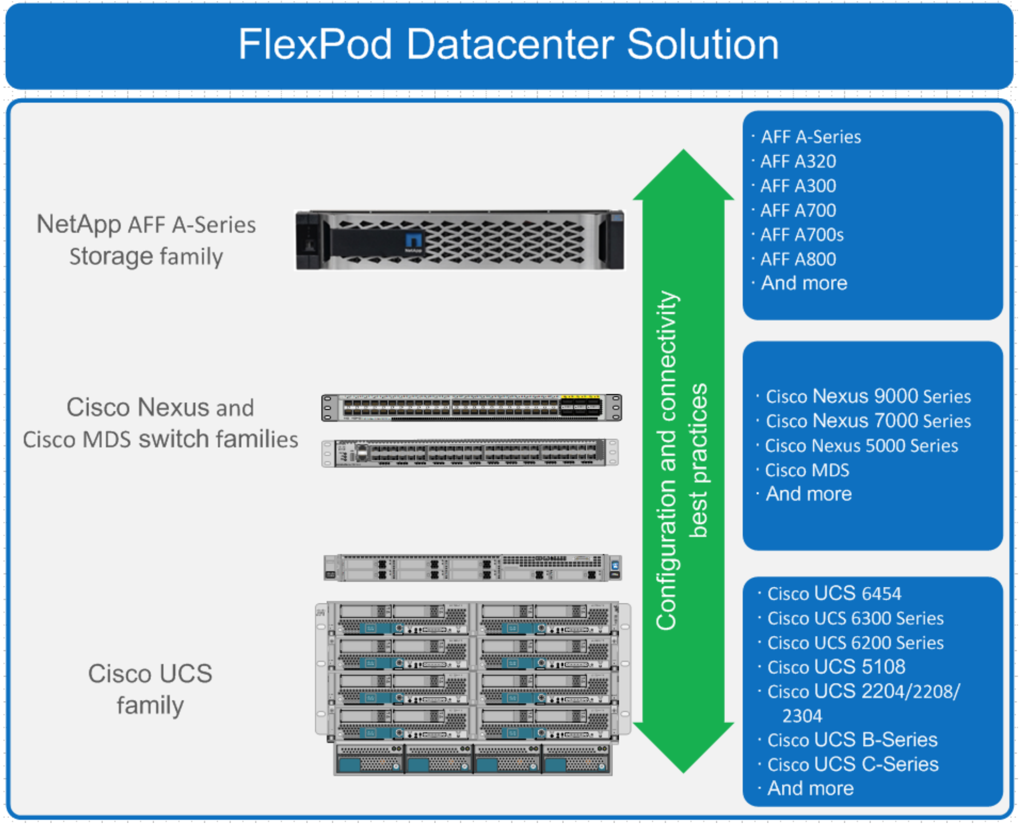
FlexPod Datacenter Solution
NetApp AFF A-Series Storage family
Cisco Nexus and Cisco MDS switch families
Cisco UCS family
Configuration and connectivity best practices
· AFF A-Series
· AFF A320
· AFF A300
· AFF A700
· AFF A700s
· AFF A800
· And more
· Cisco Nexus 9000 Series
· Cisco Nexus 7000 Series
· Cisco Nexus 5000 Series
· Cisco MDS
· And more
· Cisco UCS 6454
· Cisco UCS 6300 Series
· Cisco UCS 6200 Series
· Cisco UCS 5108
· Cisco UCS 2204/2208/2304
· Cisco UCS B-Series
· Cisco UCS C-Series
· And more

These components are connected and configured according to the combined best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence. FlexPod can be scaled up for greater performance and capacity by adding compute, network, or storage resources individually as needed. It can also be scaled out for environments that require multiple consistent deployments, such as by rolling out of additional FlexPod stacks. The reference architecture covered in this document uses the Cisco Nexus 9000 for the network switching element.

One of the key benefits of FlexPod is its ability to maintain consistency during scaling. Each of the component families shown (Cisco UCS, Cisco Nexus, and NetApp AFF) offers platform and resource options to scale the infrastructure up or down, while still supporting the same features and functionality that are required for the configuration and connectivity best practices of FlexPod.

Cisco Nexus

Cisco Nexus series switches provide an Ethernet switching fabric for communications between the Cisco UCS, NetApp storage controllers, and the rest of a customer's network. There are many factors to consider when choosing the main data switch in this type of architecture to support both the scale and the protocols required for the resulting applications. All Nexus switch models including the Nexus 5000 and Nexus 7000 are supported in this design and may provide additional features such as FCoE or OTV. However, be aware that there may be slight differences in setup and configuration based on the switch used. The validation for this deployment leverages the Cisco Nexus 9300 series switches, which deliver high performance 100/40GbE ports, density, low latency, and exceptional power efficiency in a broad range of compact form factors.

The Cisco Nexus 9000 family of switches supports two modes of operation: NxOS standalone mode and Application Centric Infrastructure (ACI) fabric mode. In standalone mode, the switch performs as a typical Nexus switch with increased port density, low latency and 40Gb connectivity. In fabric mode, the administrator can take advantage of Cisco ACI. Cisco Nexus 9000 based FlexPod design with Cisco ACI consists of Cisco Nexus 9500 and 9300 based spine/leaf switching architecture controlled using a cluster of three Application Policy Infrastructure Controllers (APICs).

Many of the most recent single-site FlexPod designs also use this switch due to the advanced feature set and the ability to support Application Centric Infrastructure (ACI) mode. When leveraging ACI fabric mode, the Nexus 9000 series switches are deployed in a spine-leaf architecture. The design guide captures both reference architectures: one with Nexus Switches in the standalone mode and another with ACI mode.

For more information, refer to http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html.

This standalone Nexus based FlexPod design deploys a single pair of Nexus 9336C-FX2 top-of-rack switches within each placement, using the traditional standalone mode running NX-OS.

The traditional deployment model delivers numerous benefits for this design:

- High performance and scalability with L2 and L3 support per port
- Layer 2 multipathing with all paths forwarding through the Virtual port-channel (vPC) technology
- VXLAN support at line rate
- Advanced reboot capabilities include hot and cold patching
- Hot-swappable power-supply units (PSUs) and fans with N+1 redundancy

Cisco Nexus 9000 provides an Ethernet switching fabric for communications between the Cisco UCS domain, the NetApp storage system and the enterprise network. In the FlexPod design, Cisco UCS Fabric Interconnects and NetApp storage systems are connected to the Cisco Nexus 9000 switches using virtual Port Channels (vPC).

Cisco Application Centric Infrastructure (ACI)

Cisco Nexus 9000 Series Switches are the foundation of the ACI architecture and provide the network fabric. A new operating system is used by Cisco Nexus 9000 switches running in ACI mode. The switches are then coupled with a centralized controller, the APIC, and its open API. The APIC is the unifying point of automation, telemetry, and management for the ACI fabric, helping to enable an application policy model approach to the data center.

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

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

The Cisco ACI fabric consists of three major components:

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

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

Figure 2 Cisco ACI Fabric Architecture

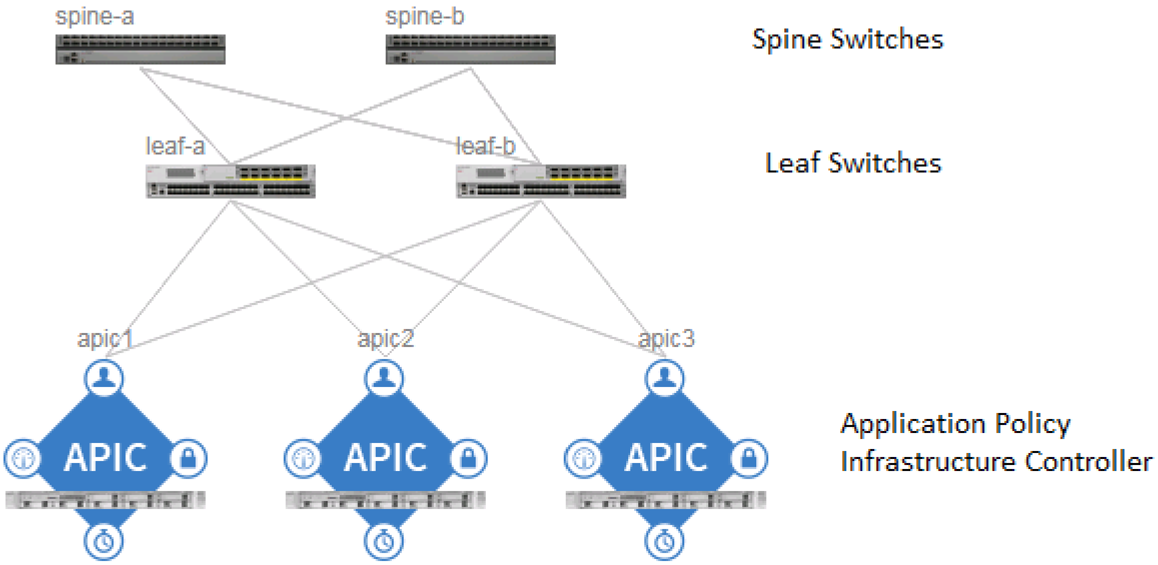

ACI Components

- Cisco Application Policy Infrastructure Controller

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

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

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

APIC is responsible for all tasks enabling traffic transport including:

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

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

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

For more information on Cisco APIC, refer to:

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

- Leaf Switches

In Cisco ACI, all workloads connect to leaf switches. Leaf switch, typically can be a fixed form Nexus 9300 series or a modular Nexus 9500 series switch that provides physical server and storage connectivity as well as enforces ACI policies. The latest Cisco ACI fixed form factor leaf nodes allow connectivity up to 25 and 40 Gbps to the server and uplinks of 100 Gbps to the spine. There are a number of leaf switch choices that differ based on functions like port speed, medium type, multicast routing support, scale of endpoints, and so on.

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

- Spine Switches

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

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

Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10-100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- Compute - The compute piece of the system incorporates servers based on the Second-Generation Intel® Xeon® Scalable processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.
- Network - The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lowers costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- Virtualization - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.
- Storage access – Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.
- Management: The system uniquely integrates compute, network and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired physical configurations and infrastructure policies for applications. Service Profiles increase business agility by enabling IT to automate and provision resources in minutes instead of days.

Cisco UCS Differentiators

Cisco's Unified Computing System is revolutionizing the way servers are managed in the datacenter. The following are the unique differentiators of Cisco Unified Computing System and Cisco UCS Manager.

- Embedded Management — In Cisco UCS, 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.
- Unified Fabric — In Cisco UCS, from blade server chassis or rack servers 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 the overall solution.

- Auto Discovery — By simply inserting the blade server in the chassis or connecting the rack server to the fabric interconnect, discovery and inventory of compute resources occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN and management networks.

- Policy Based Resource Classification — Once 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 showcases 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.

- Model based Management Architecture — The Cisco UCS Manager architecture and management database is model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.

- Policies, Pools, Templates — The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a 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 flexibility 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 another policy by name is resolved in the organizational hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then the special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.

- Service Profiles and Stateless Computing — A 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 — The combination of policies, pools and templates, loose referential integrity, policy resolution in the organizational hierarchy and a service profiles-based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environments typically observed in private and public clouds.

- Extended Memory — The enterprise-class Cisco UCS Blade server extends the capabilities of Cisco’s Unified Computing System portfolio in a half-width blade form factor. It harnesses the power of the latest Intel® Xeon® Scalable Series processor family CPUs and Intel® Optane DC Persistent Memory (DCPMM) with up to 18TB of RAM (using 256GB DDR4 DIMMs and 512GB DCPMM).

- Simplified QoS — Even though Fibre Channel and Ethernet are converged in the Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, integrated management for all software and hardware components in Cisco UCS. Using Cisco Single Connect technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive graphical user interface (GUI), a command-line interface (CLI), or a through a robust application programming interface (API).

Cisco UCS Manager is embedded into the Cisco UCS Fabric Interconnects and provides a unified management interface that integrates server, network, and storage. Cisco UCS Manger performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers comprehensive set of XML API for third party integration, exposes thousands of integration points and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Cisco UCS™ Manager, Release 4.0 provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) across multiple chassis and Cisco UCS servers. Cisco UCS Manager, Release 4.0 is a unified software release for all supported Cisco UCS hardware platforms. Release 4.0 enables support for UCS 6454 Fabric Interconnects, VIC 1400 series adapter cards on UCS M5 servers and Second Generation Intel® Xeon® Scalable processor refresh and Intel® Optane™ Data Center persistent memory modules on UCS Intel-based M5 servers.

For more information on Cisco UCSM Release 4.0 refer to the Release Notes page.

Cisco Intersight

Cisco Intersight is Cisco's new systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent management level and enables IT organizations to analyze, simplify and automate their IT environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings in Total Cost of Ownership (TCO) and to deliver applications faster to support new business initiatives.

The Cisco UCS platform uses model-based management to provision servers and fabric automatically, regardless of form factor. Cisco Intersight works in conjunction with Cisco UCS Manager and the Cisco Integrated Management Controller (IMC). By simply associating a model-based configuration with a resource through service profiles, your IT staff can consistently align policy, server personality, and workloads. These policies can be created once and used by IT staff with minimal effort to deploy servers. The result is improved productivity and compliance and lower risk of failures due to inconsistent configuration.

Cisco Intersight will be integrated with data center, hybrid cloud platforms and services to securely deploy and manage infrastructure resources across data center and edge environments. In addition, Cisco will provide future integrations to third-party operations tools to allow customers to use their existing solutions more effectively.

Cisco Intersight manages all Cisco UCS servers and switches in the solution and offers cloud-based, centralized management of Cisco UCS servers across all Enterprise locations and delivers unique capabilities such as:

- Integration with Cisco TAC for support and case management
- Proactive, actionable intelligence for issues and support based on telemetry data
- Compliance check through integration with Cisco Hardware Compatibility List (HCL)

- Centralized service profiles for policy-based configuration

For more information about Cisco Intersight and the different editions, go to: Cisco Intersight – Manage your systems anywhere.

Cisco UCS Fabric Interconnects

The Cisco UCS Fabric interconnects (FIs) provide a single point for connectivity and management for the entire Cisco UCS system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by the Cisco UCS Manager. Cisco UCS FIs provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, SAN and management traffic using a single set of cables.

The 3rd generation (6300) Fabric Interconnect deliver options for both high workload density, as well as high port count, with both supporting either Cisco UCS B-Series blade servers, or Cisco UCS C-Series rack mount servers. The Fabric Interconnect models featured in this design is Cisco UCS 6332-16UP Fabric Interconnect which is a 1RU 40GbE/FCoE switch and 1/10 Gigabit Ethernet, FCoE and FC switch offering up to 2.24 Tbps throughput. The switch has 24x40Gbps fixed Ethernet/FCoE ports with unified ports providing 16x1/10Gbps Ethernet/FCoE or 4/8/16Gbps FC ports. This model is aimed at FC storage deployments requiring high performance 16Gbps FC connectivity to Cisco MDS switches or FC direct attached storage.

Figure 3 Cisco UCS 6332-16UP Fabric Interconnect

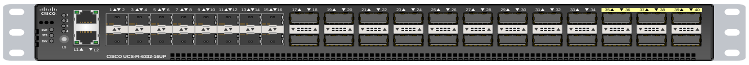

See the Solution References section for additional information on Fabric Interconnects.

Cisco UCS 2304 XP Fabric Extenders

The Cisco UCS Fabric extenders (FEX) or I/O Modules (IOMs) multiplexes and forwards all traffic from servers in a blade server chassis to a pair of Cisco UCS Fabric Interconnects over a 40Gbps unified fabric links. All traffic, including traffic between servers on the same chassis, or different chassis, is forwarded to the parent fabric interconnect where Cisco UCS Manager runs, managing the profiles and polices for the servers. FEX technology was developed by Cisco. Up to two FEXs can be deployed in a chassis.

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 4 Cisco UCS 2304 XP Fabric Extenders

See the Solution References section for additional information on Fabric Extenders.

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a fundamental building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server architecture. The Cisco UCS blade server chassis uses an innovative unified fabric with fabric-extender technology to lower TCO by reducing the number of network interface cards (NICs), host bus adapters (HBAs), switches, and cables that need to be managed, cooled, and powered. It is a 6-RU chassis that can house up to 8 x half-width or 4 x full-width Cisco UCS B-series blade servers. A passive mid-plane provides up to 80Gbps of I/O bandwidth per server slot and up to 160Gbps for two slots (full-width). The rear of the chassis contains two I/O bays to house a pair of Cisco UCS 2000 Series Fabric Extenders to enable uplink connectivity to FIs for both redundancy and bandwidth aggregation.

Figure 5 Cisco UCS 5108 Blade Server Chassis

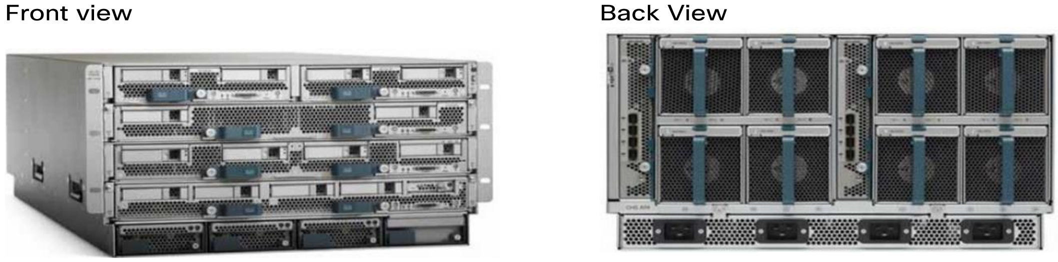

Cisco UCS 1400 Series Virtual Interface Cards (VICs)

Cisco VICs support Cisco SingleConnect technology, which provides an easy, intelligent, and efficient way to connect and manage computing in your data center. Cisco SingleConnect unifies LAN, SAN, and systems management into one simplified link for rack servers and blade servers. This technology reduces the number of network adapters, cables, and switches needed and radically simplifies the network, reducing complexity. Cisco VICs can support 256 Express (PCIe) virtual devices, either virtual Network Interface Cards (vNICs) or virtual Host Bus Adapters (vHBAs), with a high rate of I/O Operations Per Second (IOPS), support for lossless Ethernet, and 10/25/40/100-Gbps connection to servers. The PCIe Generation 3 x16 interface helps ensure optimal bandwidth to the host for network-intensive applications, with a redundant path to the fabric interconnect. Cisco VICs support NIC teaming with fabric failover for increased reliability and availability. In addition, it provides a policy-based, stateless, agile server infrastructure for your data center.

The VIC 1400 series is designed exclusively for the M5 generation of UCS B-Series Blade Servers and C-Series Rack Servers. The adapters can support 10/25/40/100-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.

Cisco UCS B200 M5 Blade Servers

The Cisco UCS B200 M5 Blade Server shown in Figure 6, is a half-width blade successor from the Cisco UCS B200 M4.

Figure 6 Cisco UCS B200 M5 Blade Server

It features:

- 2nd Generation Intel® Xeon® Scalable processors with up to 28 cores per socket
- Up to 3 terabytes (TB) of DDR4 memory for improved performance

- Up to 7.5 terabytes (TB) using 12x128G DDR4 DIMMs and 12x512G Intel® Optane DCPMM nonvolatile memory technology
- Up to two NIVIDA GPUs
- Two Small-Form-Factor (SFF) drive slots
- Up to two Secure Digital (SD) cards or M.2 SATA drives

For more information about the Cisco UCS B200 M5 Blade Servers, see the Cisco UCS B200 M5 Blade Server datasheet.

Cisco UCS B480 M5 Servers

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. The Cisco UCS B480 M5 is a full-width blade server supported by the Cisco UCS 5108 Blade Server Chassis.

The Cisco UCS B480 M5 Blade Server offers Four Intel Xeon Scalable CPUs (up to 28 cores per socket), up to 12 TB of DDR4 memory and 18 TB using 24x256G DDR4 DIMMs and 24x512G Intel® Optane DC Persistent Memory. Five mezzanine adapters and support for up to four GPUs and Cisco UCS Virtual Interface Card (VIC) 1440 modular LAN on Motherboard (mLOM) and Cisco UCS Virtual Interface Card (VIC) 1480 is a dual-port 40-Gbps Ethernet.

Figure 7 Cisco UCS B480 M5 Blade Server

NetApp AFF A320 Storage

With the new NetApp® AFF A-Series controller lineup, NetApp provides industry leading performance combined with a full suite of enterprise-grade data-management and data-protection features.

This architecture uses the NetApp AFF A320 all-flash array as the foundation for infrastructure storage. The AFF A320 controllers provides the high-performance benefits of 100GbE and NVMe all-flash solid-state drives (SSDs) by using the external NetApp NS224 storage shelf. The shelf is connected through the NVMe/RoCE protocol and offers 24 bays for NVMe SSDs.

High-speed network connectivity is achieved with 100GbE Ethernet and 32Gb FC fabric connectivity. Multiple expansion slots per controller can support 10GbE, 25GbE, 40GbE, and 100GbE Ethernet connectivity as well as 8Gb, 16Gb, and 32Gb FC networking.

The ability to scale out to 24 storage nodes makes the AFF A320 the ideal storage controller for shared workload needs within a converged infrastructure.

For more information about the NetApp AFF A320 and all the AFF A-series controllers, see the AFF product page.

You can view or download more technical specifications of the AFF A-series controllers using the [NetApp datasheet ds-3582](#).

Figure 8 NetApp AFF A320 and NS224 Storage Shelf

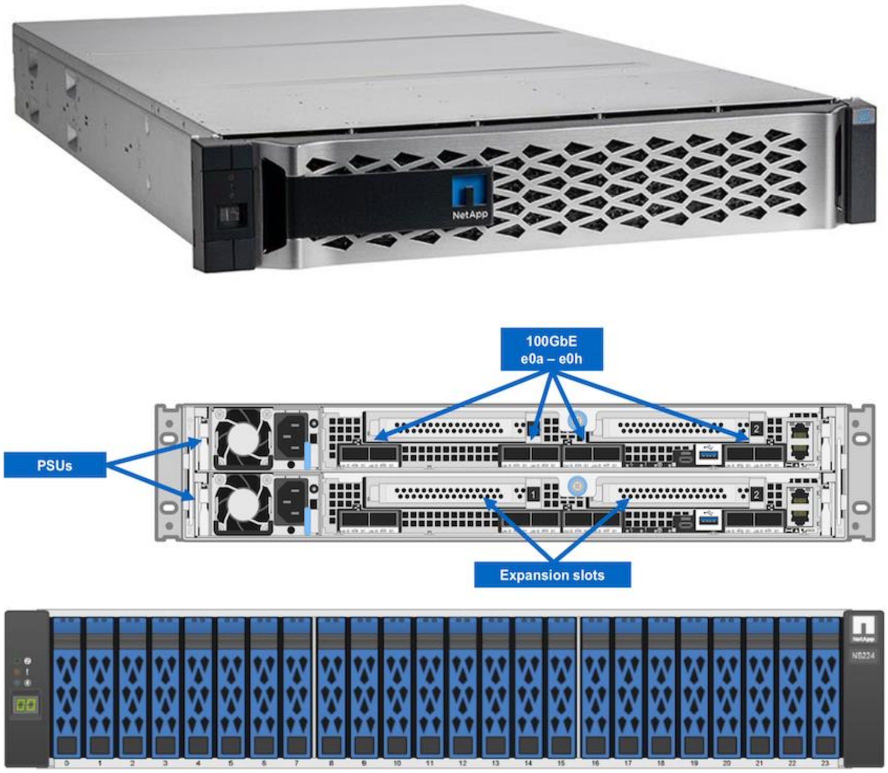

NetApp ONTAP 9.6

NetApp ONTAP 9.6 data management software is used with the NetApp AFF A320 all-flash storage system in this solution design. ONTAP software offers unified storage for applications that read and write data over block or file-access protocol storage configurations.

ONTAP implementations can run on NetApp engineered FAS or AFF series arrays. They can also run on commodity hardware (NetApp ONTAP Select) and in private, public, or hybrid clouds (NetApp Cloud Volumes ONTAP and the NetApp Cloud Volumes Service). Specialized implementations offer best-in-class converged infrastructure, featured here as part of the FlexPod Datacenter solution, or with access to third-party storage arrays (NetApp FlexArray® virtualization).

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

The following sections provide an overview of ONTAP 9.6 as an industry-leading data management software architected on the principles of software-defined storage.

For more information about all the capabilities of ONTAP data management software, see: [ONTAP Data Management Software](#).

NetApp Storage Virtual Machine

A NetApp ONTAP cluster serves data through at least one, and potentially multiple, storage virtual machines (SVMs). An SVM is a logical abstraction that represents the set of physical resources of the cluster. SVMs enable multitenancy with strict separation of data for each tenant. Data volumes and network LIFs are created and assigned to an SVM and can reside on any node in the cluster that the SVM has access to.

An SVM can own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node in the storage cluster to another. For example, a NetApp FlexVol® flexible volume can be non-disruptively moved to a new node, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and therefore it is not tied to any specific piece of physical hardware.

An SVM can support multiple data protocols concurrently, and the volumes within the SVM can be joined to form a single NAS namespace. The namespace makes all the SVM's data available through a single share or mount point to NFS and CIFS clients. SVMs also support block-based protocols, and LUNs can be created and exported by using iSCSI, FC, or FCoE. Any or all of these data protocols can be used within a given SVM. Storage-administrator and storage-management roles can be associated with an SVM, offering higher security and access control. This security is important in environments that have more than one SVM and when storage is configured to provide services to different groups or sets of workloads.

Storage Efficiencies

Storage efficiency is a primary architectural design element of ONTAP data management software. A wide array of features enables you to store more data while using less space. In addition to deduplication and compression, you can store your data more efficiently by using features such as unified storage, multitenancy, thin provisioning, and by leveraging NetApp Snapshot™ technology.

Compaction (introduced in ONTAP 9) is the latest patented storage efficiency technology released by NetApp. In the NetApp WAFL® file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on disk to save space. ONTAP 9.6 includes enhancements in the compression and compaction algorithms. These new storage efficiency technologies allow ONTAP to store more data in less space, reducing storage costs and maximizing the effective capacity of your storage system.

Encryption

Data security remains an important consideration for customers purchasing storage systems. Before ONTAP 9, NetApp supported full-disk encryption in storage clusters. In ONTAP 9, however, the encryption capabilities of ONTAP were extended with an Onboard Key Manager (OKM). The OKM generates and stores keys for each of the drives in ONTAP, enabling ONTAP to provide all of the functionality required for encryption directly out of the box. Through this functionality, known as NetApp Storage Encryption (NSE), sensitive data stored on disk is secure and can only be accessed by the ONTAP storage system with the correct keys.

NetApp has extended the encryption capabilities of ONTAP further with NetApp Volume Encryption (NVE), a software-based mechanism for encrypting data. NVE allows you to encrypt data at the per-volume level instead of requiring the encryption of all data within the cluster, providing more flexibility and granularity to ONTAP administrators. This encryption extends to Snapshot copies and NetApp FlexClone® volumes that are created within the cluster.

One benefit of NVE is that it runs after the implementation of the storage efficiency features, and, therefore, it does not interfere with the ability of ONTAP to create space savings. NVE unifies data encryption available on-premises and extends it out into the cloud. NVE is also FIPS 140-2 compliant. This compliance helps businesses adhere to federal regulatory guidelines for data stored within the cloud. Aggregate-level encryption is new in ONTAP 9.6. This functionality offers aggregate-wide deduplication, which was not previously available with NVE.

Other enhancements of ONTAP 9.6 are the key management support at the SVM level, self-encrypting drives, default cluster peering encryption, and the support of wire encryption for NetApp SnapMirror® technology.

For more information about encryption in ONTAP, see the NetApp Power Encryption Guide in the NetApp ONTAP 9 Documentation Center.

FlexClone

NetApp FlexClone technology allows you to create nearly instantaneous point-in-time copies of a FlexVol volume without consuming any additional storage until the cloned data changes from the original. FlexClone volumes add extra agility and efficiency to storage operations. It takes only a few seconds to create a FlexClone volume and doing so does not interrupt access to the parent FlexVol volume. FlexClone volumes use space efficiently, applying the ONTAP architecture to store only data that changes between the parent and clone. FlexClone volumes are suitable for testing or development environments or for any environment in which progress is made by locking-in incremental improvements. FlexClone volumes also improve any business process in which you must distribute data in a changeable form without endangering the integrity of the original.

SnapMirror (Data Replication)

NetApp SnapMirror® is an asynchronous replication technology for data replication across different sites, within the same data center, between an on-premises datacenter and the cloud, or for a cloud to on-premises datacenter. ONTAP 9.5 introduces volume granular, zero-data-loss protection with SnapMirror Synchronous (SM-S). SnapMirror Synchronous extends traditional SnapMirror volume replication to synchronous mode, so that you can meet zero recovery-point-objective (RPO) disaster-recovery and compliance objectives. Policy-based replication provides a simple and familiar configuration interface that is managed with the same tools as traditional SnapMirror, including the ONTAP CLI, ONTAP System Manager, Active IQ Unified Manager, the NetApp Manageability SDK, and ONTAP RESTful APIs.

In addition, over-the-wire encryption for NetApp SnapMirror technology has been available since ONTAP 9.6, increasing security for data replication.

NetApp SnapCenter

NetApp SnapCenter® is next-generation data protection software for tier-1 enterprise applications. SnapCenter, with its single-pane-of-glass management interface, automates and simplifies the manual, complex, and time-consuming processes associated with the backup, recovery, and cloning of multiple databases and other application workloads.

SnapCenter uses software such as NetApp Snapshot copies, SnapMirror replication, NetApp SnapRestore® data recovery, and FlexClone thin cloning. These features enable it to integrate seamlessly with technology from Oracle, Microsoft, SAP, VMware, and MongoDB across the FC, iSCSI, and NAS protocols. This integration enables IT organizations to scale their storage infrastructure, meet increasingly stringent SLA commitments, and improve the productivity of administrators across the enterprise. SnapCenter is used in this solution for backup and restore of SAP HANA systems.

SnapCenter Architecture

SnapCenter is a centrally managed, web-based application that runs on a Windows platform and manages and protects multiple servers remotely.

Figure 9 illustrates the high-level architecture of a NetApp SnapCenter Server.

Figure 9 SnapCenter Architecture

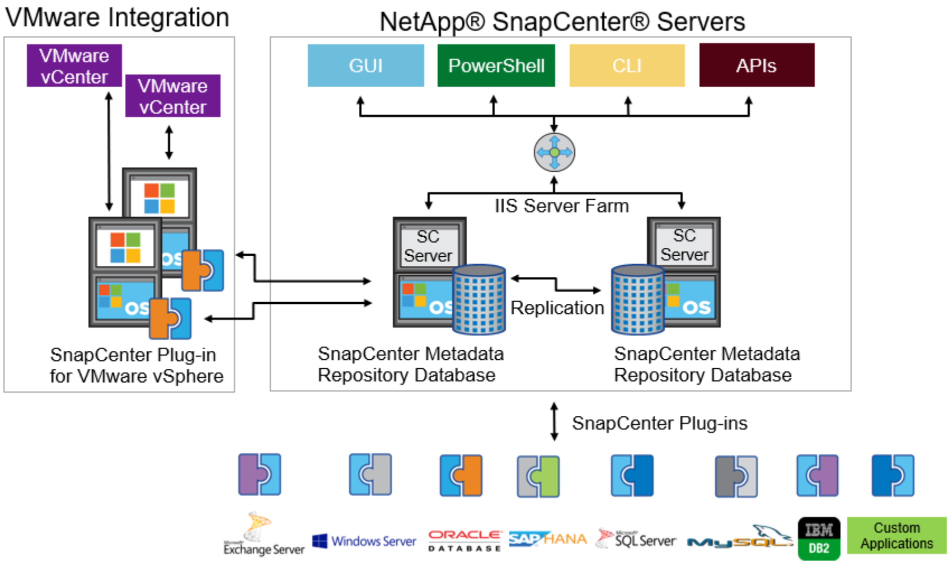

The SnapCenter Server has an HTML5-based GUI and Windows PowerShell cmdlets and APIs. The SnapCenter Server is capable of high availability out of the box. If one SnapCenter host becomes unavailable for any reason, then the second SnapCenter Server can take over seamlessly and no operations are affected.

The SnapCenter Server can push out plug-ins to remote hosts. These plug-ins are used to interact with applications, databases, or file systems. Usually, the plug-ins must be present on the remote host so that application-level or database-level commands can be issued from the same host where the application or database is running.

SnapCenter uses SM Service to manage plug-ins and interactions between the SnapCenter Server and the plug-in host. SM Service is a NetApp SnapManager® web service running on top of Windows Server internet information services (IIS) on SnapCenter Server. SM Service takes all client requests such as backup, restore, and clone.

The SnapCenter Server communicates those requests to SMCore, a service that runs within the SnapCenter Server and remote servers. SMCore plays a significant role in coordinating with the SnapCenter plug-ins package for Windows.

SnapCenter Virtualization (SCV) is a plug-in that manages virtual servers running on VMware and helps to discover the host file system, databases on virtual machine disks (VMDKs), and raw device mapping (RDM).

SnapCenter Features

SnapCenter enables you to create application-consistent Snapshot copies and to complete data protection operations, including Snapshot copy-based backup, clone, restore, and backup-verification operations. SnapCenter provides a centralized management environment, and it uses role-based access control (RBAC) to delegate data protection and management functions to individual application users across the SnapCenter Server and Windows hosts.

SnapCenter includes the following key features:

- A unified and scalable platform across applications and database environments with virtual and nonvirtual storage powered by the SnapCenter Server

- Consistency of features and procedures across plug-ins and environments supported by the SnapCenter UI
- Role-based access control (RBAC) for security and centralized role delegation
- Application-consistent Snapshot copy management, restore, clone, and backup verification support from both primary and secondary destinations (NetApp SnapMirror and NetApp SnapVault® technology)
- Remote package installation from the SnapCenter GUI
- Nondisruptive, remote upgrades
- A dedicated SnapCenter repository for faster data retrieval
- Load balancing that is implemented by using Microsoft Windows network load balancing (NLB) and application request routing (ARR) with support for horizontal scaling
- Centralized scheduling and policy management to support backup and clone operations
- Centralized reporting, monitoring, and dashboard views
- SnapCenter 4.2 support for data protection for VMware virtual machines, SQL Server databases, Oracle databases, MySQL, SAP HANA, MongoDB, and Microsoft Exchange

SAP HANA Data Protection with SnapCenter

The FlexPod solution can be extended with additional software and hardware components to cover data protection, backup and recovery, and disaster recovery operations. The following chapter provides a high-level overview of how to enhance SAP HANA backup and disaster recovery using the NetApp SnapCenter plug-In for SAP HANA.

More details on the setup and configuration of SnapCenter for backup and recovery or disaster recovery operations can be found in the following technical reports:

- SAP HANA Backup and Recovery with SnapCenter
- SAP HANA Disaster Recovery with Asynchronous Storage Replication

SAP HANA Backup

Storage-based Snapshot backups are a fully supported and integrated backup method available for SAP HANA.

Storage-based Snapshot backups are implemented with the NetApp SnapCenter plug-in for SAP HANA, which creates consistent Snapshot backups by using the interfaces provided by the SAP HANA database. SnapCenter registers the Snapshot backups in the SAP HANA backup catalog so that they are visible within the SAP HANA studio or cockpit and can be selected for restore and recovery operations.

Snapshot copies created within primary storage can be replicated to the secondary backup storage by using NetApp SnapMirror technology controlled by SnapCenter. Different backup retention policies can be defined for backups held on the primary storage and to those backups held on the secondary storage. The SnapCenter Plug-In for SAP HANA manages the retention of Snapshot copy-based data backups and log backups, including housekeeping of the backup catalog. The SnapCenter plug-in for SAP HANA also allows the execution of a block integrity check of the SAP HANA database by executing a file-based backup.

Storage-based Snapshot backups provide significant advantages when compared to file-based backups. Advantages include the following:

- Rapid backup (less than a minute)
- Faster restore on the storage layer (less than a minute)
- No performance effect on the SAP HANA database host, network, or storage during backup
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

SAP HANA Disaster Recovery with Asynchronous Storage Replication

SAP HANA disaster recovery can be performed either on the database layer by using SAP system replication or on the storage layer by using storage replication technologies. This section provides an overview of disaster recovery solutions based on asynchronous storage replication.

The same SnapCenter plug-in that is described in the section “SAP HANA Backup” is also used for the asynchronous mirroring solution. A consistent Snapshot image of the database at the primary site is asynchronously replicated to the disaster recovery site with SnapMirror.

High-level Architecture Description

Figure 10 shows a high-level overview of the data protection architecture.

For an offsite backup and disaster recovery solution, the following additional hardware and software components are required:

- A Windows host to run SnapCenter server software
- Offsite backup storage to replicate backups from primary storage to a secondary storage system
- Disaster recovery storage to replicate backups from primary storage to a disaster recovery site

The SnapCenter server must be able to communicate with the SVMs that are used at the primary (within the FlexPod instance), the offsite backup location, and the disaster recovery storage.

The primary storage must have a network connection to the offsite storage and the disaster recovery storage. A storage cluster peering must be established between the primary storage, the offsite storage, and the disaster recovery storage.

The SnapCenter server must have a network connection to the SAP HANA database hosts to deploy the HANA plug-in and to communicate with the plug-in after deployment. As an alternative, the HANA plug-in can also be deployed at the FlexPod management server. See SAP HANA Backup and Recovery with SnapCenter for more details on the deployment options for the HANA plug-in.

Figure 10 Data Protection with SnapCenter

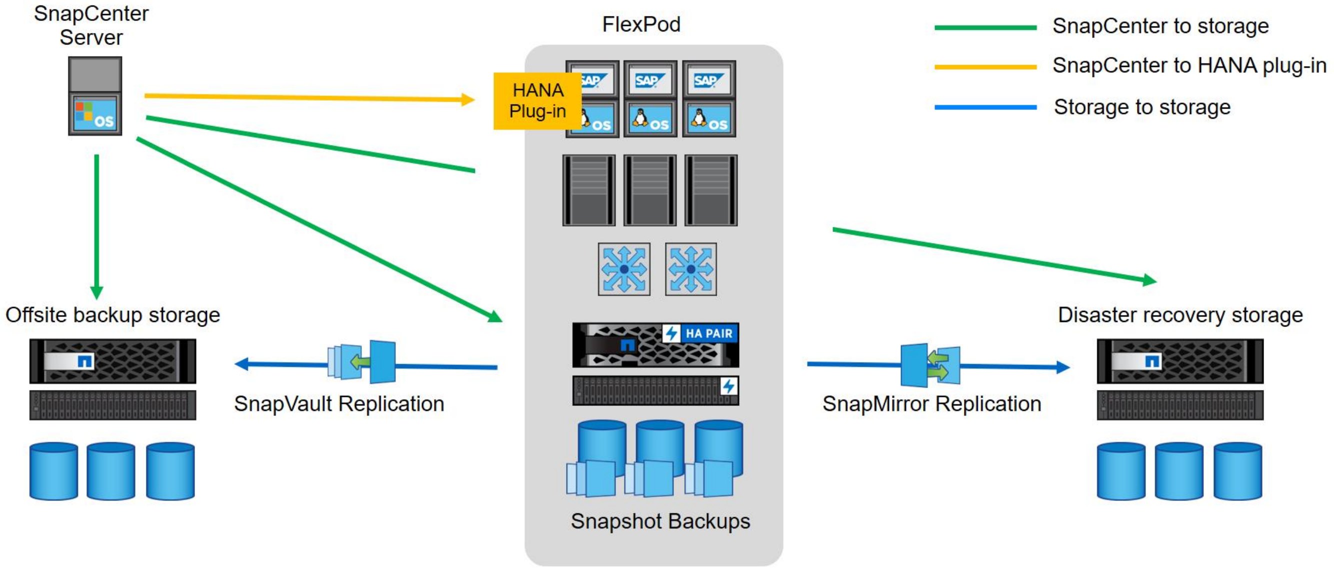

SAP Landscape Management

SAP Landscape Management (LaMa) enables SAP system administrators to automate SAP system operations, including end-to-end SAP system copy and refresh operations. SAP LaMa is one of the few SAP software products with which infrastructure providers such as NetApp and Cisco can integrate their products. With such integration, customers can employ NetApp functions directly from the SAP LaMa GUI.

NetApp offers NetApp Storage Services Connector (SSC) that allows SAP LaMa to directly access technologies and features such as NetApp FlexClone® instant cloning and NetApp SnapMirror data replication. These technologies help minimize storage use and shorten the time required to create SAP system clones and copies.

With the help of the built-in functions and a rich set of extensibility features within SAP LaMa, FlexPod customers can directly integrate storage-based backups or instantaneously create space-efficient FlexClone system copies on the primary datacenter. They can even use storage at either the offsite backup or disaster recovery site.

Figure 11 shows how SAP LaMa and NetApp SSC can be integrated into the overall FlexPod architecture.

Figure 11 SAP Landscape Management

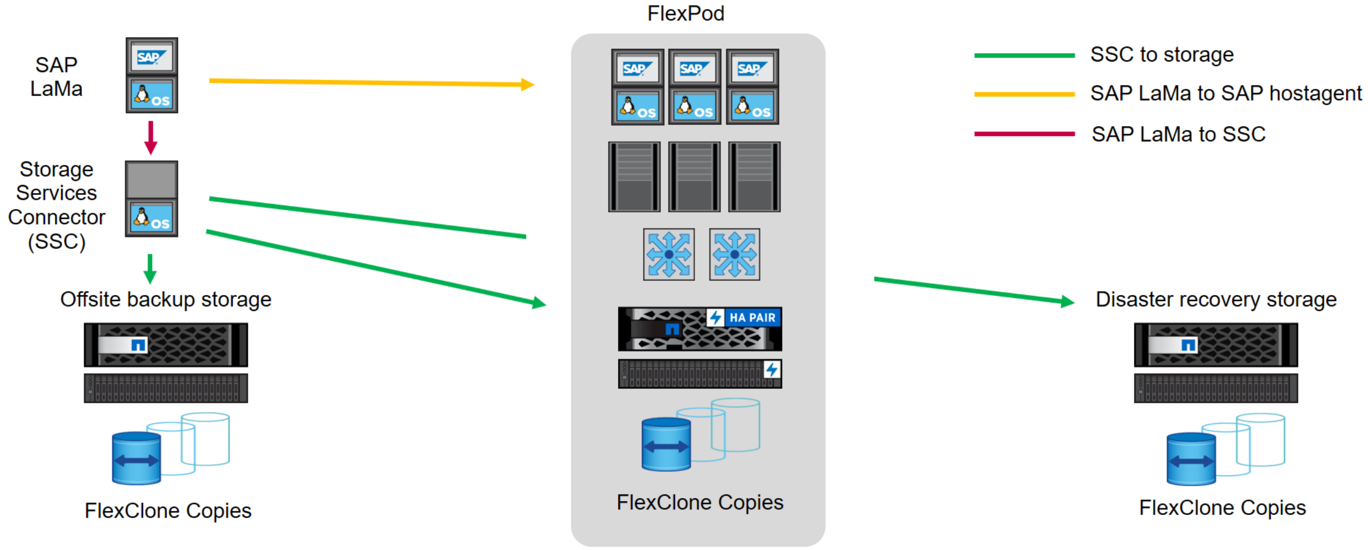

From an administrator's perspective, SAP LaMa is the central tool to operate and monitor SAP systems, compute instances, and required storage resources. Figure 11 illustrates the required network communications between the different components.

- SAP LaMa must be able to communicate with SAP Host Agent running on the physical or virtual host. Although SAP Host Agent is automatically installed during an SAP system installation, it can be manually configured to include hosts in SAP LaMa management that do not run SAP software, such as web servers.
- To communicate with NetApp storage systems, SAP LaMa must be able to communicate with NetApp SSC. For more information about NetApp SSC, see the NetApp SSC for SAP LaMa site.
- NetApp SSC version 4.0 is an executable that must be installed onto a Linux host that is accessible by SAP LaMa and is also able to connect to all NetApp storage systems integrated into SAP LaMa.

For a detailed description of SAP LaMa and the NetApp Storage Services Connector, see the technical report Integrating NetApp ONTAP systems with SAP Landscape Management.

The SAP LaMa does not include regular backup and recovery or disaster recovery functionality. These functionalities are provided by SnapCenter.

SAP Application Monitoring with AppDynamics

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

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

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

For more information, go to: https://docs.appdynamics.com/display/SAP/SAP+Monitoring+Using+AppDynamics

Solution Design

The SAP HANA TDI option enables multiple SAP HANA production systems to run on the same infrastructure. In this configuration, the existing blade servers used by different SAP HANA systems share the same network infrastructure and storage systems. In addition, the SAP application server can share the same infrastructure as the SAP HANA database.

SAP HANA in FlexPod environment qualify as SAP HANA TDI implementations.

The section explains the SAP HANA system requirements defined by SAP and Architecture of FlexPod Datacenter Solution providing the platform for SAP and SAP HANA.

SAP HANA System Implementation Options

Because multiple implementation options are available specific to this shared TDI usage, you need to define your requirements before you can select the solution components. This section defines the basic requirements for each option.

Single SAP HANA System on a Single Server: Scale-Up (Bare Metal or Virtualized)

A scale-up TDI solution is the simplest of the installation types. In general, this solution provides the best SAP HANA performance. All data and processes are located on the same server and can be accessed locally, and no network communication with other SAP HANA nodes is required. SAP HANA scale-up TDI solutions are based on a standalone rack-mount server or blade server and use the intended external storage.

The network requirements for this option depend on the client and application access and storage connections. If you don’t need system replication or a backup network, a 1 Gigabit Ethernet (access) network and bandwidth factored for the data, log and shared filesystem access storage networks are required to run SAP HANA in a scale-up configuration.

The storage IO as well as latency Key Performance Indicators (KPI) requirements by SAP need to be fulfilled.

For a virtualized SAP HANA installation, remember that hardware components, such as network interfaces and host bus adapters (HBAs), are shared. If possible, use dedicated network adapters and HBAs for each virtualized SAP HANA system. The storage requirements are similar to a bare metal use case with respect to storage IO and latency KPIs.

Single SAP HANA System on Multiple Servers: Scale-Out

You should use a scale-out TDI solution if the SAP HANA system does not fit into the main memory of a single server based on the rules defined by SAP. With 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.

The network requirements for this option are higher than for scale-up systems. In addition to the client and application access and storage access networks, you also must have a node-to-node network. If you don’t need system replication or a backup network, a 10 Gigabit Ethernet (access) network and a mandatory minimum of 10

Gigabit Ethernet (node-to-node) and bandwidth factored for the data, log and shared filesystem access storage networks are required to run SAP HANA in a scale-out configuration.

The storage configuration should make sure that the SAP requirements for IO and the latency KPI are fulfilled.

For a virtualized SAP HANA installation, remember that hardware components, such as network interfaces and host bus adapters (HBAs), are shared. If possible, use dedicated network adapters and HBAs for each virtualized SAP HANA system. The network and storage requirements are similar to a bare metal use case with respect to storage IO/latency and inter-node network bandwidth.

Multiple SAP HANA Systems: Scale-Up (Bare Metal or Virtualized)

Using the SAP HANA TDI option for shared storage and a shared network, you can build a solution to run multiple SAP HANA scale-up systems on shared infrastructure. One approach is to use a SAP HANA scale-out solution and install one SAP HANA system per server. A 4 + 1 scale-out solution (four active nodes and one standby node) includes all the components needed to run five SAP HANA systems based on the TDI key performance indicators (KPIs).

The network requirements are the same as for a single SAP system.

The storage IO as well as latency KPI requirements by SAP need to be fulfilled for each individual system at all times.

For a virtualized SAP HANA installation, remember that some hardware components are shared: for example, network interfaces and HBAs. If possible, you should use dedicated network adapters and HBAs for each virtualized SAP HANA system. The storage requirements are similar to a bare metal use case with respect to storage IO and latency KPIs.

Multiple SAP HANA Systems: Scale-Out (Bare Metal or Virtualized)

Using the SAP HANA TDI option for shared storage and a shared network, you can build a solution to run multiple SAP HANA scale-out systems on shared infrastructure. One approach is to use a SAP HANA scale-out solution infrastructure and install two or more SAP HANA systems on it. As an example, you can use an 11 + 1 scale-out solution (11 active nodes and 1 standby node) includes all the components to install two 5 + 1 systems or three 3 + 1 systems or any other supported scale-out configuration.

The network requirements are the same as for a single SAP HANA system.

The storage IO as well as latency KPI requirements by SAP need to be fulfilled for each individual scale-out system at all times.

For a virtualized SAP HANA installation, remember that some hardware components are shared: for example, network interfaces and HBAs. If possible, you should use dedicated network adapters and HBAs for each virtualized SAP HANA system. The network and storage requirements are similar to a bare metal use case with respect to storage IO/latency and inter-node network bandwidth.

Co-existing SAP HANA and SAP Application Workloads

With SAP HANA TDI it is possible to run SAP HANA on shared infrastructure that also hosts non HANA workloads as standard SAP applications. Scenarios where SAP HANA database bare metal installation along with virtualized SAP application workloads are common in the datacenter. It is important make sure there is appropriate storage IO and network bandwidth segregation so that HANA systems get their due to comfortably satisfy the storage and network KPIs for production support.

Scaling-Up and Scaling-Out SAP HANA Systems

Hosting multiple scale-up and scale-out systems call for proper sizing of the infrastructure with a clear compute node to storage system ratio. The number of compute nodes along with storage arrays has to be determined based on the number total number of SAP HANA nodes that would make up the system landscape and would involve corresponding scaling of associated compute gear and networking components based on port availability and usage.

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: SAP HANA Hardware Directory

CPU

With the release of the Second-Generation Intel® Xeon® Scalable processors (Cascade Lake), SAP supports Intel Xeon Platinum CPUs with 28 cores per CPU in SAP HANA environments. The Cisco UCS B-Series Blade Servers are capable to be configured with full or half size amount of Intel Xeon scalable family CPUs.

Memory

The Cisco Integrated Management Controller (IMC) and Cisco UCS Manager Release 4.0(4) introduce support for Intel® Optane™ Data Center persistent memory modules (DCPMM) on Cisco UCS M5 servers based on the Second-Generation Intel ® Xeon® Scalable processors (Cascade Lake).

Detailed information on the configuration and management is available in the whitepaper Cisco UCS: Configuring and Managing Intel Optane Data Center Persistent Memory Modules.

Intel Optane DCPMMs must be installed with DRAM dual in-line memory modules (DIMM) in the same system and will not function without any DRAM DIMMs installed. In two- or four-socket configurations, each socket contains two IMCs. Each memory controller connects to three double data rate (DDR) memory channels which connects to two physical DIMM/persistent memory slots.

SAP HANA 2.0 SPS03 revision 35+ currently support various memory capacity ratios between Intel Optane DCPMM and DRAM DIMMs in the same system and will not function without any DRAM DIMMs installed.

In DDR4 DIMM memory only population the following configuration rules apply:

- 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
- Supported Memory Configuration for SAP NetWeaver Business Warehouse (BW) and DataMart
  - 1.5 TB on Cisco UCS B200 M5 Servers with 2 CPUs
  - 3 TB on Cisco UCS B480 M5 Servers with 4 CPUs
- Supported Memory Configuration for SAP Business Suite on SAP HANA (SoH)
  - 3 TB on Cisco UCS B200 M5 Servers with 2 CPUs

  - 6 TB on Cisco UCS B480 M5 Servers with 4 CPUs

In Intel Optane DCPPM/DDR4 DIMM mixed memory population the following rules apply:

- Maximum use of all available memory channels
- Supported Memory Configuration for SAP Business Suite on SAP HANA (SoH)
  - 7.5 TB on Cisco UCS B200 M5 Servers with 2 CPUs
  - 18 TB on Cisco UCS B480 M5 Servers with 4 CPUs

Network

An 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. Different clients, such as SQL clients on SAP application servers, browser applications using HTTP/S to the SAP HANA XS server and other data sources (such as BI) need a network communication channel to the SAP HANA database.
- Internal zone. The internal zone covers the communication between hosts in a distributed SAP HANA system as well as the communication used by SAP HANA system replication between two SAP HANA sites.
- Storage zone. Although SAP HANA holds the bulk of its data in memory, the data is also saved in persistent storage locations. In most cases, the preferred storage solution involves separate, externally attached storage subsystem devices that can provide dynamic mount-points for the different hosts, according to the overall landscape. A storage area network (SAN) can also be used for storage connectivity.

Storage

SAP HANA is an in-memory database which uses storage devices to save a persistent 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 a storage system in the SAP HANA TDI option, the storage must be certified as SAP HANA certified Enterprise Storage.

The Solution References section provides links to the SAP HANA certified hardware directory and a white paper which discuss all relevant information about the storage requirements.

The solution needs to pass the SAP HANA KPI check successfully prior of reporting IO performance related SAP HANA incidents.

File System Layout

Figure 12 illustrates the SAP HANA file system layout and the recommended storage sizes to install and operate SAP HANA. The recommendation is to reserve for the Linux operating system root volume 10GB of disk space and to store the SAP software 50GB of disk space. In this solution the root volume /root and SAP software /usr/sap are in the same disk volume, although they can be setup in two different volumes as well.

Figure 12 File System Layout for 2 Node Scale-Out System

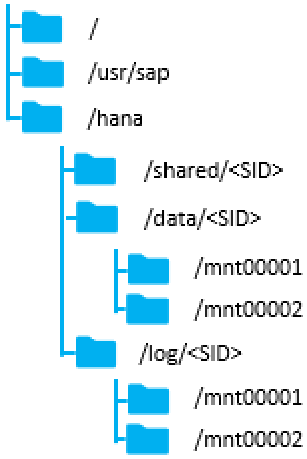

The sizing for SAP HANA file system volumes is based on the amount of memory equipped on the SAP HANA host.

Scale-Up Solutions

The recommended, minimum disk space requirements for SAP HANA TDI installations are:

/ (root)	100 GB inclusive of space required for /usr/sap
/hana/shared	1 × RAM or 1TB whichever is less
/hana/data	1 × RAM
/hana/log	512 GB

Scale-Out Solutions

The recommended, minimum disk space requirements for SAP HANA TDI installations are:

/ (root)	100 GB inclusive of space required for /usr/sap
/hana/shared	1 × RAM for every 4 active HANA nodes
/hana/data	1 × RAM for each active HANA node
/hana/log	512 GB for each active HANA node

Operating System

The supported operating systems for SAP HANA with Intel® Optane™ DCPMM are, as follows:

- SUSE Linux Enterprise Server for SAP Applications 15 GA
- Red Hat Enterprise Linux for SAP HANA 7.6

High Availability

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

- 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 integrated high-availability functionality. If an 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.

Although not tested and validated in this design, additional high-availability solutions like SAP HANA System Replication with Linux Cluster are available as well. For detailed information, refer to the SAP HANA Administration Guide – High Availability for SAP HANA or SAP HANA Administration Guide – Configuring SAP HANA System Replication.

Physical Topology

Figure 13 shows the FlexPod Datacenter solution components for SAP HANA and the network connections for a configuration with the Cisco UCS 6300 Fabric Interconnects. This design has port-channeled 40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects and 40Gb /100 Gb Ethernet connections between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000s, and between Cisco Nexus 9000s and NetApp AFF A320 storage array.

Figure 13 FlexPod Topology

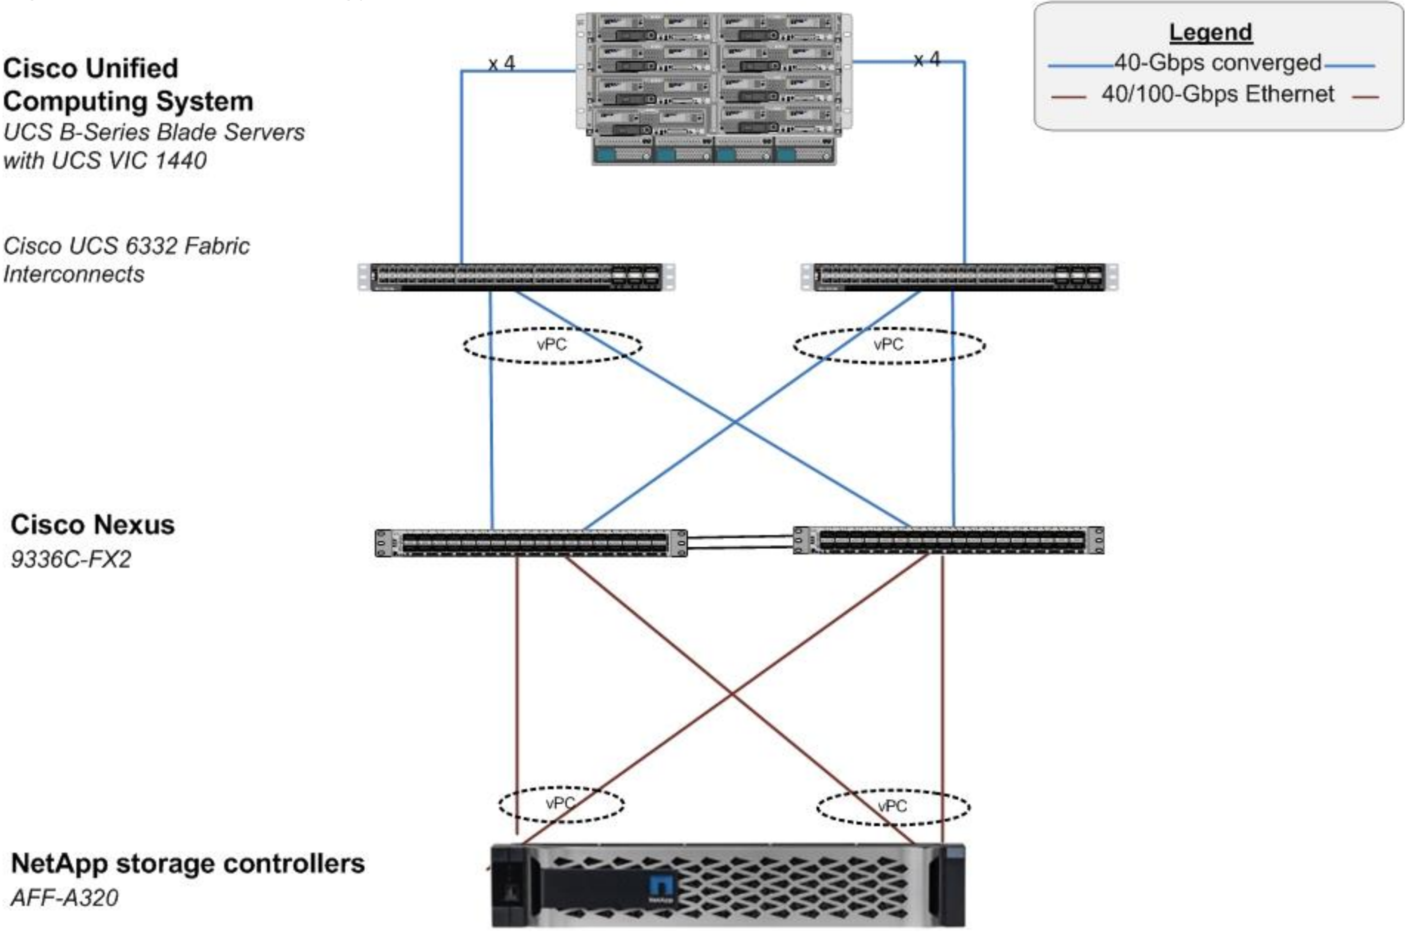
Cisco Unified Computing System
UCS B-Series Blade Servers with UCS VIC 1440
x 4
x 4
Legend
40-Gbps converged
40/100-Gbps Ethernet
Cisco UCS 6332 Fabric Interconnects
vPC
vPC
Cisco Nexus
9336C-FX2
vPC
vPC
NetApp storage controllers
AFF-A320

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

The reference 100Gb based hardware configuration includes:

- Two Cisco Nexus 9336c-FX2 switches
- Two Cisco UCS 6300 series fabric interconnects
- One NetApp AFF A320 (HA pair) running ONTAP 9.6 with one NS224 NVMe SSD disk shelf

Compute Connectivity

Each compute chassis in the design is connected to the managing fabric interconnect with at least two ports per IOM. These connections from the Cisco UCS 6332-16UP Fabric Interconnect to the 2304 IOM are shown in Figure 14.

Figure 14 Cisco UCS 6332-16UP Fabric Interconnect Connectivity to FEX 2304 IOM on Chassis 5108

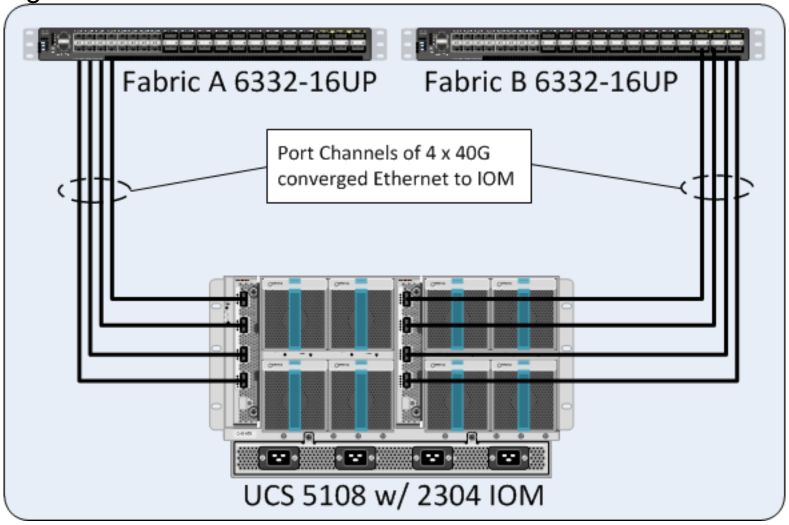

The 2304 IOM are shown with 4 x 40Gbps ports to deliver an aggregate of 320 Gbps to the chassis.

Network Connectivity

Nexus switches form the access layer. Both compute via Fabric Interconnects and NetApp storage controllers connect to Nexus 9336C-FX2 pair.

The network coming into each of the fabric interconnects is configured as a Port Channel to the respective fabric interconnects but is implemented as Virtual Port Channels (vPC) from the upstream Nexus switches. The same applies for the storage controllers' connection. In this design, the 40GE ports were used for the construction of the port channels that connected to the vPCs (Figure 15).

Figure 15 Cisco UCS 6332-16UP Fabric Interconnect Uplink Connectivity to Nexus 9336-FX2

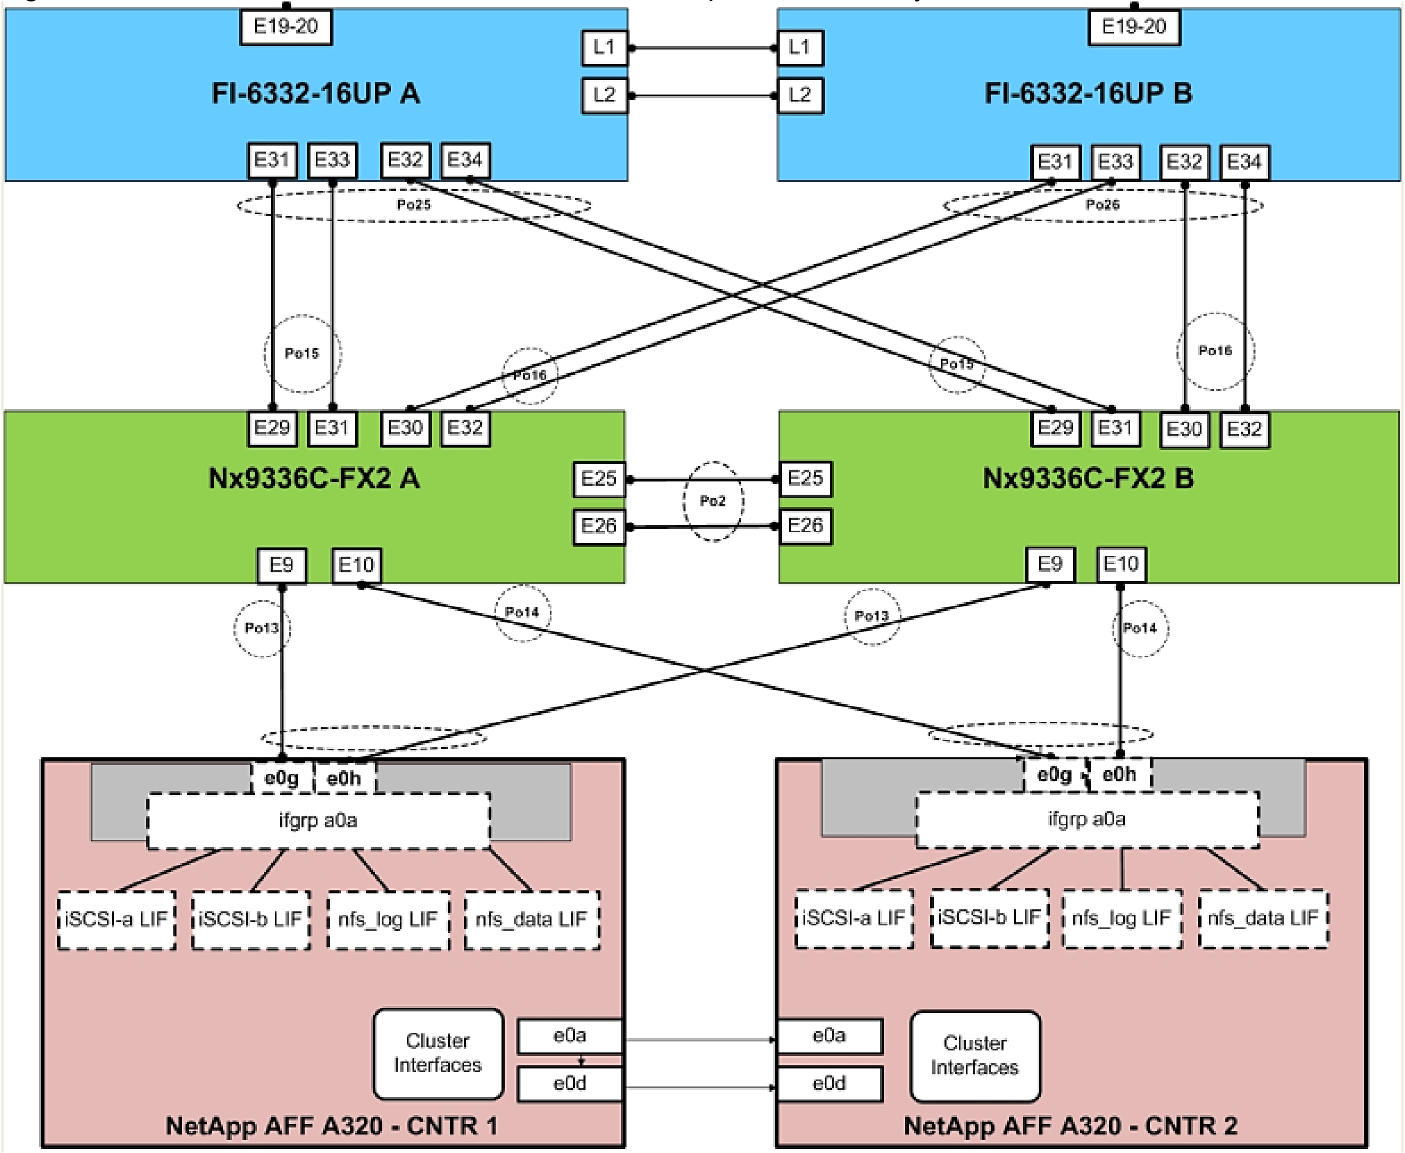

Figure 15 shows the connections between the Cisco Nexus 9000s, Cisco UCS Fabric Interconnects and NetApp AFF A320s. A vPC requires a "peer link" which is documented as port channel 2 in this diagram. In addition to the vPC peer-link, the vPC peer keepalive link is a required component of a vPC configuration. The peer keepalive link allows each vPC enabled switch to monitor the health of its peer. This link accelerates convergence and reduces the occurrence of split-brain scenarios. In this validated solution, the vPC peer keepalive link uses the out-of-band management network.

vPC Considerations

- Define a unique domain ID
- Set the priority of the intended vPC primary switch lower than the secondary (default priority is 32768)
- Establish peer keepalive connectivity. It is recommended to use the out-of-band management network (mgmt0) or a dedicated switched virtual interface (SVI)
- Enable vPC auto-recovery feature

- Enable peer-gateway. Peer-gateway allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer allowing vPC peers to forward traffic
- Enable IP ARP synchronization to optimize convergence across the vPC peer link.
- A minimum of two 10 Gigabit Ethernet connections are required for vPC
- All port channels should be configured in LACP active mode

Spanning Tree Considerations

- The spanning tree priority was not modified. Peer-switch (part of vPC configuration) is enabled which allows both switches to act as root for the VLANs
- Loopguard is disabled by default
- BPDU guard and filtering are enabled by default
- Bridge assurance is only enabled on the vPC Peer Link.
- Ports facing the NetApp storage controller and Cisco UCS are defined as "edge" trunk ports

For configuration details, refer to the Cisco Nexus 9000 Series Switches Configuration guides: http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-installation-and-configuration-guides-list.html.

Network Bandwidth Considerations

It is recommended to go with a architecture design that allows to use a min of 25GE single stream bandwidth and upwards from compute through storage. Reference architecture with 3rd Gen FIs enabling 40GE end to end even while the upstream Nexus Switches and NetApp AFF 320 array allow for 100GE connect between the them addresses the SAP HANA inter-node network bandwidth requirement with comfortable ease.

Storage Connectivity

The storage controllers' 100GbE onboard ports for data networking are directly connected to the Cisco Nexus 9336C-FX2 switch ports in a port-channel configuration. Figure 16 shows the port and interface assignment connection diagram for the AFF storage to the Cisco Nexus 9336C-FX2 network fabrics. In this design, NFS and iSCSI traffic uses the 40GbE bandwidth.

Figure 16 Storage Connectivity Overview

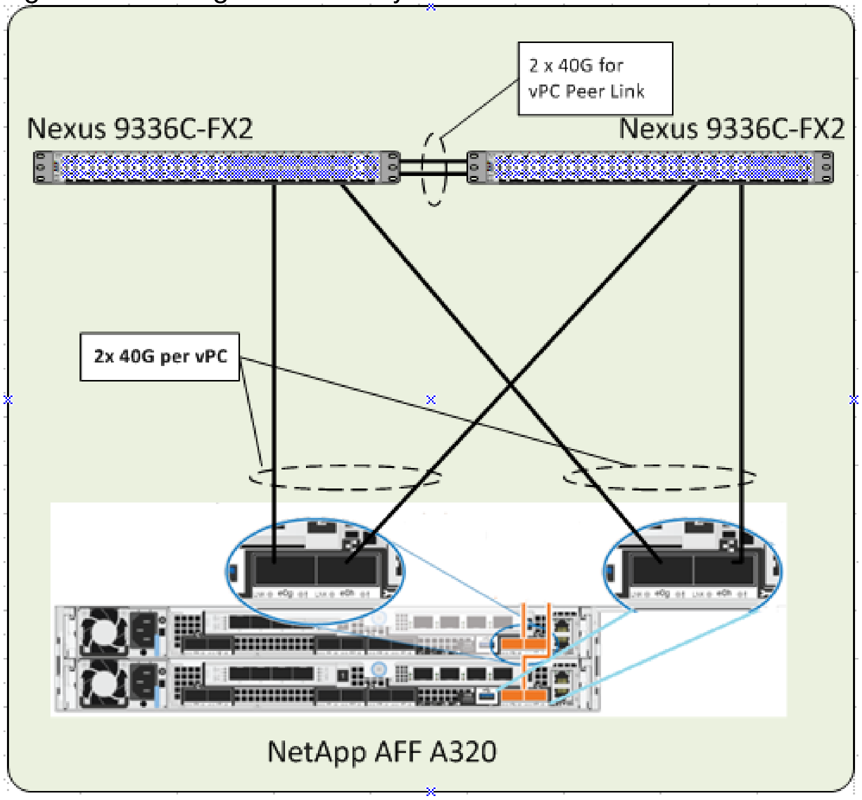

End-to-End IP Network Connectivity

The Cisco Nexus 9000 is the key component that brings together the 40Gbps capabilities of the other parts of this design. vPCs extend to both the AFF A320 controllers and the Cisco UCS 6332 fabric interconnects. Passage of this traffic (shown in Figure 17 from left to right) is as follows:

- Coming from the Cisco UCS B480 M5 server, equipped with a VIC 1440/1480 adapter, allowing for 40Gb using two ports out of four of IOM on each side of the fabric (A/B) into the server
- Pathing through 40Gb KR lanes of the Cisco UCS 5108 Chassis backplane into the Cisco UCS 2304 IOM (fabric extender)
- Connecting from each IOM to the fabric interconnect with up to four 40Gb links automatically configured as port channels during chassis association
- Continuing from the Cisco UCS 6332 fabric interconnects into the Cisco Nexus 9336C-FX2 with a bundle of 40Gb ports presenting each side of the fabric from the Nexus pair as a common switch using a vPC
- Ending at the AFF A320 controllers with 40/100GbE bundled vPCs from the Nexus switches now carrying both sides of the fabric.

Figure 17 vPC, AFF A320 Controllers, and Cisco UCS 6332 Fabric Interconnect Traffic

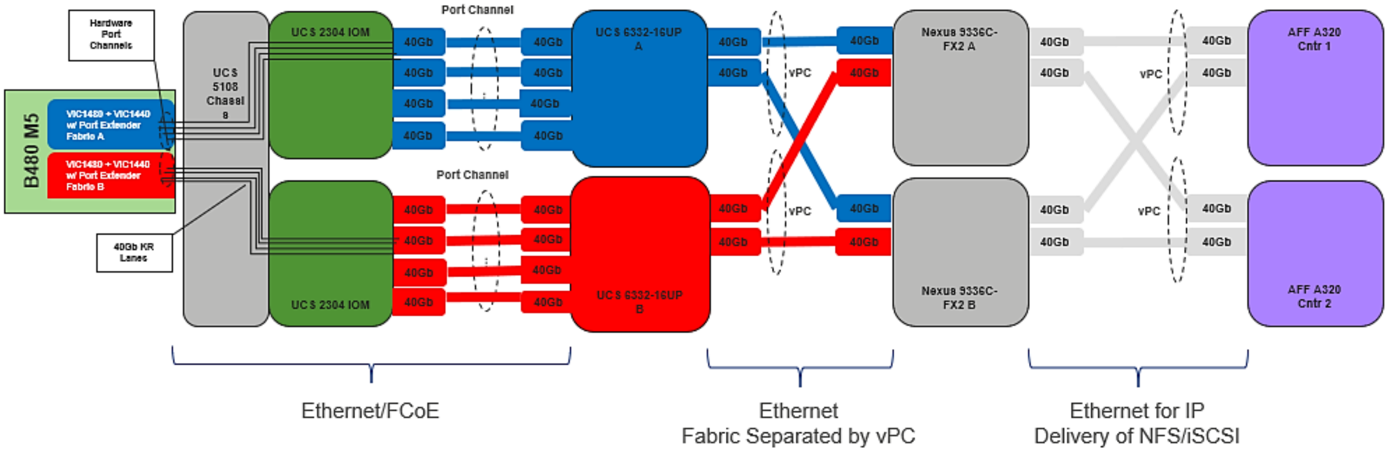

The equivalent view for a Cisco UCS C-Series server is shown in Figure 18. The main difference is that the two 40Gbps interfaces are connected directly between the VIC 1495/1497 and the fabric interconnect and are port-channeled into an 80Gbps interface:

Figure 18 Cisco UCS C-Series Server

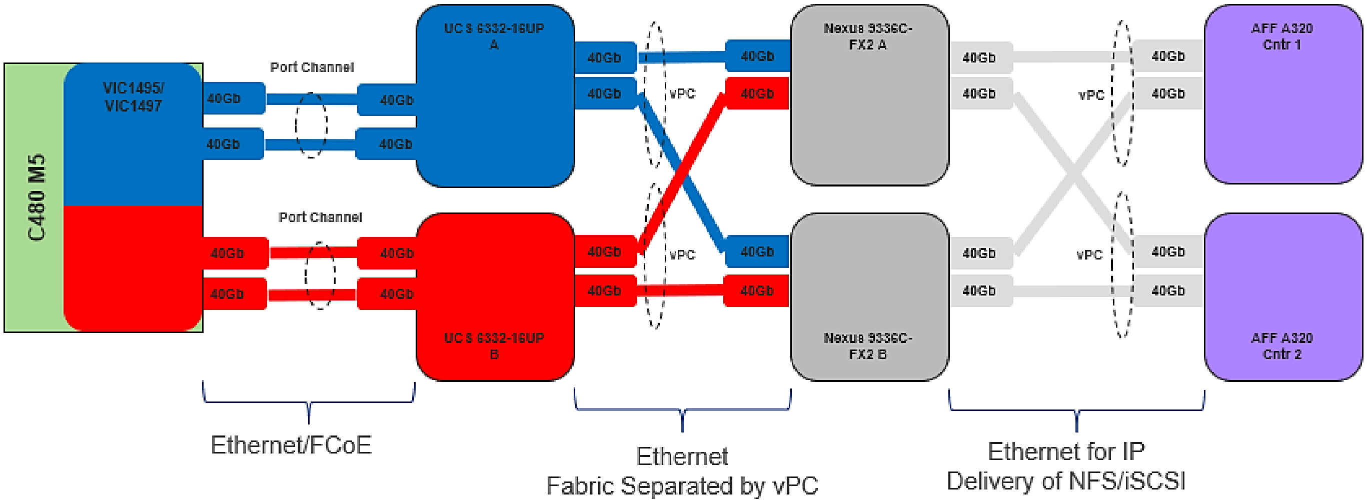

Compute Design Options

Cisco UCS B-Series

The Cisco UCS B-Series servers are configured in the design with:

- iSCSI boot – Persistent operating system installation, independent of the physical blade for true stateless computing.
- VIC 1440 with Port Expander and VIC 1480 provides four 40Gbps capable of up to 256 Express (PCIe) virtual adapters.

SAP HANA supports Intel® Optane™ Data Center Persistent Memory Module (DCPMM). Intel Optane DCPMM is supported on Cisco Unified Computing System™ (Cisco UCS®) servers for SAP HANA.

Memory for databases is currently small, expensive, and volatile. Intel Optane DC persistent memory is denser, more affordable, and persistent, and it performs at speeds close to that of memory. These features of Intel Optane DC persistent memory can help lower TCO through reduced downtime and simplified data-tiering operations. These same features can also make SAP HANA in-memory databases economically viable for a wider range of use cases. Intel Optane DC persistent memory provides near-DRAM in-memory computing speed in a form factor similar to that of dual in-line memory modules (DIMMs) at a lower price per gigabyte than DRAM. . With its persistence, performance, and lower cost per gigabyte than conventional memory, Intel Optane DC persistent memory can help reduce total cost of ownership (TCO), reshape the way that businesses tier their data for database systems, and open new use cases for the speed and power of the SAP HANA platform.

The FlexPod Datacenter solution for SAP HANA could be based on 4 socket Cisco UCS B480 M5 Blade Server and C480 M5 Rack Server or 2 socket Cisco UCS B200 M5 Blade Server and C220 and C240 M5 Rack Servers

Table 1 and Table 2 summarizes the server specifications with possible memory configurations for the SAP HANA use case, respectively both with DRAM and

Table 1 Overview of Cisco UCS B480 M5 Blade and Cisco UCS C480 M5 Rack Server Configuration

CPU specifications	Intel Xeon Platinum 8276L/8280L processor: Quantity 4
Possible memory configurations	32-GB DDR4: Quantity 24 (768 GB) 64-GB DDR4: Quantity 24 (1.5 TB) 128-GB DDR4: Quantity 24 (3 TB)
Possible DCPMM memory configurations	128-GB DCPMM: Quantity 24 (3 TB) 256-GB DCPMM: Quantity 24 (6 TB) 512-GB DCPMM: Quantity 24 (12 TB)

Table 2 Overview of Cisco UCS C240, Cisco UCS C220 M5 Rack Server, and Cisco UCS B200 M5 Blade Server Configuration

CPU specifications	Intel Xeon Platinum 8276L/8280L processor: Quantity 2
Possible memory configurations	16-GB DDR4: Quantity 12 (192 GB) 32-GB DDR4: Quantity 12 (384 GB) 64-GB DDR4: Quantity 12 (768 TB) 128-GB DDR4: Quantity 12 (1.5 TB)
Possible DCPMM memory configurations	128-GB DCPMM: Quantity 12(1.5 TB) 256-GB DCPMM: Quantity 12 (3 TB) 512-GB DCPMM: Quantity 12 (6 TB)

Intel Optane DCPMMs must be installed with DRAM DIMMs in the same system. The persistent memory modules will not function without any DRAM DIMMs installed. In two-, four-, and eight-socket configurations, each socket contains two IMCs. Each memory controller is connected to three double data rate (DDR) memory channels that are then connected to two physical DIMM persistent memory slots.

SAP HANA 2.0 SPS 03 currently supports various capacity ratios between Intel Optane DCPMMs and DIMMs.

For information regarding the Cisco UCS compute with Intel Optane DC Persistent Memory Module (DCPMM) and possible capacity ratios between DCPMMs and DIMMs, refer to: https://www.cisco.com/c/dam/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/whitepaper-c11-742627.pdf

Cisco UCS M5 servers with second-generation Intel Xeon Scalable processors and Intel Optane DC persistent memory, combined with DRAM, revolutionizes the SAP HANA landscape by helping organizations achieve lower overall TCO, ensure business continuity, and increase the memory capacities of their SAP HANA deployments.

Network Design Options

Management Connectivity

Out-of-band management is handled by an independent switch that could be one currently in place in the customer's environment. Each physical device had its management interface carried through this Out-of-band switch, with in-band management carried as a different VLAN within the solution for SAP HANA Network requirements.

Out-of-band configuration for the components configured as in-band could be enabled however would require additional uplink ports on the 6332-16UP Fabric Interconnects if the out of band management is kept on a separate out of band switch. A disjoint layer-2 configuration can then be used to keep the management and data plane networks separate. This would require additional vNICs on each server, which are then associated with the management uplink ports.

Jumbo Frames

Jumbo frames are a standard recommendation across Cisco designs to help leverage the increased bandwidth availability of modern networks. To take advantage of the bandwidth optimization and reduced consumption of CPU resources gained through jumbo frames, they were configured at each network level to include the virtual switch and virtual NIC.

This optimization is relevant for VLANs that stay within the pod, and do not connect externally. Any VLANs that are extended outside of the pod should be left at the standard 1500 MTU to prevent drops from any connections or devices not configured to support a larger MTU.

Storage Design Options

NetApp AFF Storage System Family

As all NetApp AFF storage systems use ONTAP as the storage operating system, the functionality of ONTAP is available starting with entry class systems, over mid-range systems, and all the way up to high end systems. The design described here based on an AFF A320 can also be achieved with other NetApp AFF storage systems as long as the same ethernet connection speed of 40GBE is available for the system.
This allows customer to choose the right storage system for their needs.

Storage Design Considerations

The following NetApp storage design best practices and recommendations were used in this design:
TR-4435: SAP HANA on NetApp All Flash FAS Systems with NFS Configuration Guide
https://www.netapp.com/us/media/tr-4435.pdf

Solution Validation

A high-level summary of the FlexPod Datacenter Design validation is provided in this section. Installation procedure for both, SUSE and Red Hat Linux, following best practices from Cisco, NetApp and SAP. All SAP HANA TDI phase 5 requirements are tested and passed for performance and high availability, including:

- Cisco UCS Setup and Configuration
- NetApp Setup and Configuration
- iSCSI boot option
- Operating System Configuration for SAP HANA
- Installation of SAP HANA 2.0 SPS4
- Performance Tests using SAP's test tool

Validated Hardware and Software

Table 3 lists the hardware and software versions used during solution validation. It is important to note that Cisco and NetApp have interoperability matrixes that should be referenced to determine support for any specific implementation of FlexPod. Click the following links for more information:

- NetApp Interoperability Matrix Tool
- Cisco UCS Hardware and Software Interoperability Tool

Table 3 Validated Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6332-16UP, UCS B480 M5	UCSM 4.0(4d)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, Cisco UCS VIC 1440 and Cisco UCS VIC 1480. Also valid for rack form factor C480 M5 and two socket UCS blade and rack servers.
Network	Cisco Nexus 9336C-FX2 NX-OS	7.0(3)I7(6)	
Storage	NetApp AFF A320 with NS224 NVMe shelf	ONTAP 9.6 with NFS v4.1	
Operating System		SLES for SAP Applications 15 GA RHEL for SAP HANA 7.6	iSCSI boot

Summary

FlexPod features the latest Cisco UCS servers, Nexus fabric switches and NetApp All Flash storage. Cisco UCS has its own policy-based management with service profiles and self-integrating components. And it works with the Cisco Nexus series switches for a Unified Fabric, virtualization awareness, simple scaling, and high-performance I/O. NetApp AFF Arrays, along with ONTAP data management software defines truly unified storage that delivers the built-in storage efficiencies, integrated data protection, and intelligent management and automation.

FlexPod Datacenter is the optimal infrastructure foundation to deploy SAP HANA be it bare metal or virtualized to implement a TDI environment. It is validated for both SUSE Linux Enterprise Server and Red Hat Enterprise Linux operating systems. The solution is built utilizing Cisco UCS Blade Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches and NetApp AFF arrays. It is designed and validated using compute, network and storage best practices for high performance, scalability, and resiliency throughout the architecture. The flexibility and scalability of FlexPod also enables customers to start with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements.

References

Products and Solutions

Cisco Unified Computing System:

http://www.cisco.com/en/US/products/ps10265/index.html

Cisco UCS 6300 Fabric Interconnect:

Cisco UCS 6300 Series Fabric Interconnects

Cisco UCS 5100 Series Blade Server Chassis:

http://www.cisco.com/en/US/products/ps10279/index.html

Cisco UCS B-Series Blade Servers:

http://www.cisco.com/en/US/partner/products/ps10280/index.html

Cisco UCS C-Series Rack Mount Servers:

http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html

Cisco UCS Adapters:

http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

http://www.cisco.com/en/US/products/ps10281/index.html

Cisco Nexus 9000 Series Switches:

http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html

Cisco MDS 9000 Multilayer Fabric Switches:

https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html

AppDynamics:

https://docs.appdynamics.com/display/SAP/SAP+Monitoring+Using+AppDynamics

https://docs.appdynamics.com/display/SAP/SAP+HANA+Dashboards

NetApp ONTAP 9:

http://www.netapp.com/us/products/platform-os/ontap/index.aspx

NetApp AFF A320:

http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx

NetApp Data Management Software:

http://www.netapp.com/us/products/management-software/

NetApp SnapCenter:

https://www.netapp.com/us/products/backup-recovery/snapcenter-backup-management.aspx

TR-4435: SAP HANA on NetApp All Flash FAS Systems with NFS Configuration Guide
https://www.netapp.com/us/media/tr-4435.pdf

SAP HANA Backup and Recovery with SnapCenter

https://www.netapp.com/us/media/tr-4614.pdf

SAP HANA Disaster Recovery with Asynchronous Storage Replication

https://www.netapp.com/us/media/tr-4646.pdf

Integrating NetApp ONTAP systems with SAP Landscape Management

https://www.netapp.com/us/media/tr-4018.pdf

Interoperability Matrixes

Cisco UCS Hardware Compatibility Matrix:

https://ucshcltool.cloudapps.cisco.com/public/

NetApp Interoperability Matrix Tool:

http://support.netapp.com/matrix/

About the Authors

Pramod Ramamurthy, Technical Marketing Engineer, Cisco Systems, Inc.

Pramod is a Technical Marketing Engineer with Cisco UCS Solutions and Performance Group. Pramod has more than 15 years of experience in the IT industry focusing on SAP technologies. Pramod is currently focusing on the Converged Infrastructure Solutions design, validation and associated collaterals build for SAP HANA.

Marco Schoen, Technical Marketing Engineer, NetApp, Inc.

Marco is a Technical Marketing Engineer with NetApp and has over 16 years of experience in the IT industry focusing on SAP technologies. His specialization areas include SAP NetWeaver Basis technology and SAP HANA. He is currently focusing on the SAP HANA infrastructure design, validation and certification on NetApp Storage solutions and products including various server technologies.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Ulrich Kleidon, Cisco Systems Inc.
- Shailendra Mruthunjaya, Cisco Systems Inc.
- Sarath Gonugunta, Cisco Systems, Inc.