

VersaStack for IBM Cloud Object Storage on Cisco UCS C240 for Concentrated Dispersal Mode

Deployment Guide for IBM Cloud Object Storage on Cisco
UCS C240 M5 for Small to Medium Businesses

Last Updated: January 15, 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, see:

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

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND 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)

© 2019 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	6
Solution Overview	7
Introduction.....	7
Audience	7
Purpose of this Document.....	7
Solution Summary	8
Technology Overview	9
Cisco Unified Computing System.....	9
Cisco UCS C240 Rack Server	9
Cisco UCS Virtual Interface Card 1387	11
Cisco UCS 6300 Series Fabric Interconnect	12
Cisco Nexus 9332PQ Switch	13
Cisco UCS Manager	13
IBM Cloud Object Storage.....	15
Concentrated Dispersal Mode	16
Cloud Object Storage Components	17
Object-based Access Methods	18
REST API Access to Storage	19
Embedded Accesser.....	19
Network.....	20
IBM Hardware Requirements	20
Solution Design	22
Solution Overview.....	22
IBM Cloud Object Configuration for Cisco Validated Design	23
General Hardware Requirements	23
Compute Layer Design.....	23
Cisco UCS Server Connectivity to Unified Fabric	23
Software Distributions and Versions	26
Deployment Hardware and Software	28
Fabric Configuration	28
Configure Cisco Nexus C9332PQ Switch A and B	28
Initial Setup of Cisco Nexus C9332PQ Switch A and B	28
Enable Features on Cisco Nexus 9332PQ Switch A and B	31
Configuring VLANs on Nexus 9332PQ Switch A and B.....	32
Configuring vPC Domain on Nexus 9332PQ Switch A and B	33
Configuring Network Interfaces for vPC Peer Links on Nexus 9332PQ Switch A and B	34

Configuring Network Interfaces to Cisco UCS FI 6332 on Nexus 9332PQ Switch A and B	35
Verification Check of Cisco Nexus C9332PQ Configuration for Switch A and B	38
Initial Setup of Cisco UCS 6332 Fabric Interconnects	42
Configure Fabric Interconnect A	42
Configure Fabric Interconnect B	44
Logging Into Cisco UCS Manager	46
Configure NTP Server	46
Initial Base Setup of the Environment	47
Configure Global Policies	47
Enable Fabric Interconnect A Ports for Server	48
Enable Fabric Interconnect A Ports for Uplinks	49
Create Port Channel for Fabric Interconnect A/B	49
Label Each Server for Identification.....	50
Create KVM IP Pool.....	50
Create MAC Pool	51
Create UUID Pool.....	52
Enable CDP	53
QoS System Class	54
vNIC Template Setup	55
Adapter Policy Setup	56
Boot Policy Setup.....	57
Create Maintenance Policy Setup.....	58
Create Power Control Policy Setup	59
Create Disk Scrub Policy.....	60
Create Host Firmware Package	61
Create vMedia Policy in Cisco UCS Manager	62
Creating Storage Profiles	63
Creating Disk Group Policy for Boot Devices	63
Create Storage Profile	64
Create Service Profile Template.....	65
Create Service Profile Template	65
Identify Service Profile Template.....	65
Storage Provisioning	65
Networking.....	65
vMedia Policy	66
Server Boot Order.....	66
Server Maintenance	66

Server Assignment.....	66
Operational Policies.....	66
Create Service Profiles from Template	66
Associate Service Profiles.....	67
Installation of IBM Cloud Object Storage	68
Deployment of Virtual IBM COS Manager on VMware vCenter.....	68
Deployment of IBM COS Slicestor on Cisco UCS C240 M5L	76
IBM COS Jumbo Frame Verification	82
IBM COS dsNet Setup.....	84
Configure IBM COS to Sync with an NTP Server.....	86
Configure Access Key Authentication	87
Configure IBM COS Provisioning API	88
Create a Storage Pool	88
Create Vault for User Access.....	89
Create Vault Template for Provisioning.....	90
Create an Access Pool	91
Create a User for Object Access.....	92
Functional Object Storage Access Validation.....	94
Initial Performance S3 Benchmark with COSBench.....	95
Validation	97
IBM COS High Availability Testing.....	97
Cisco Nexus C9332PQ High Availability Testing.....	97
Cisco UCS Fabric Interconnect 6332 High Availability Testing	98
Cisco UCS C240 M5L Disk Failure Testing	98
Cisco UCS C240 M5L Node Failure Testing	99
Summary	100
About the Authors.....	101
Acknowledgements.....	101

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.

The purpose of this document is to describe the design of IBM Cloud Object Storage (COS) on latest generation of Cisco UCS C240 Rack Servers. This validated deployment provides the framework of deploying IBM COS software on Cisco UCS C240 Rack Servers. Cisco Unified Computing System (Cisco UCS) provides the storage, network, and storage access components for IBM COS, deployed as a single cohesive system.

Cisco Validated design describes how Cisco Unified Computing System can be used in conjunction with IBM COS 3.13.6 or later releases in a Concentrated Dispersal Mode. With the continuous evolution of Software Defined Storage (SDS), there has been increased demand for IBM COS solutions validated on Cisco UCS servers that can start small and grow as needed. The Cisco UCS C240 Rack Server, originally designed for the data center, together with IBM COS is ideal for such object storage solutions, making it an excellent fit for unstructured data workloads such as active archive and backup. The Cisco UCS C240 Rack Server delivers a complete infrastructure with exceptional scalability for computing and storage resources together with 25/40 Gigabit Ethernet networking.

Cisco and IBM are collaborating to offer customers a scalable object storage solution for unstructured data. Combining the power of Cisco UCS's unified management and fabric infrastructure with the reduced resource requirements of IBM COS in a Concentrated Dispersal Mode the solution is cost effective to deploy, simple to manage. This solution enables the next-generation of hybrid cloud object storage deployments driving business agility, operational efficiency and lower capital investment.

Solution Overview

Introduction

Traditional storage systems are limited in their ability to easily and cost-effectively scale to support large amounts of unstructured data. With about 80 percent of data being unstructured, new approaches using x86 servers are proving to be more cost effective, providing storage that can be expanded as easily as your data grows. Software Defined Storage is a scalable and cost-effective approach for handling large amounts of data.

However, more and more there are requirements to store unstructured data even in smaller quantities as object storage. The advantage of identifying the data by metadata and not taking over management of the location is very attractive even for smaller capacities. As a result, new technologies need to be developed to provide similar levels of availability and reliability as large scale-out object storage solutions.

IBM's COS is a storage platform that is ideal for holding large amounts of colder production data, such as backups and archives, and very large individual files, such as video files, image files, and genomic data and can also include support of warm or even hot data, by increasing CPU performance and/or memory capacity. IBM COS is highly reliable, durable, and resilient object storage that is designed for scale and security.

With the Concentrated Dispersal mode (CD mode) feature, IBM has developed a way to map the COS solution to small environments without sacrificing availability and reliability. CD mode enables customers to initially deploy IBM COS at smaller capacity sizes, starting as low as 72 terabytes, and grow incrementally adding interfaces and capacity as needed. Often customers want to start their initial investment into object storage at a small scale based on one use case. The new CD mode feature provides a cost-efficient method to begin with a small deployment (72TB) and scale to petabytes and beyond. Customers can take advantage of the benefits of IBM COS in terms of data availability and system reliability along with management capabilities, allowing them to consider extending the infrastructure to other use cases. Customers will have the flexibility to deploy low-capacity configurations in one or two-site configurations as well as with multiple sites depending on their requirements.

Together with Cisco UCS, IBM COS delivers a fully enterprise-ready solution that can manage different workloads and remain flexible. The Cisco UCS C240 Rack Server is an excellent platform to use with object workloads such as, but not limited to, active archive or backup workloads. This solution is best suited for sequential access, as opposed to random, unstructured data regardless of the data size. The Cisco UCS C240 rack server is designed to support object storage applications such as IBM COS.

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 who want to deploy an IBM COS on the Cisco Unified Computing System using Cisco UCS C240 Rack Servers.

Purpose of this Document

This document describes the architecture and deployment of an IBM COS solution using three Cisco UCS C240 Rack Servers and two Cisco UCS 6332 Fabric Interconnect managed by Cisco UCS Manager. Provided are the steps required to deploy an IBM COS in CD mode on Cisco UCS. It shows the simplicity of installing and configuring Cisco UCS rack server and illustrates the need of a well-conceived network architecture for low-latency, high-bandwidth.

Solution Summary

This Cisco Validated Design is a simple and linearly scalable architecture that provides Software Defined Storage for object on IBM COS 3.13.16 and Cisco UCS C240 Rack Server. This CVD describes in detail the design of IBM COS 3.13.16 on Cisco UCS C240 Rack Server. The solution includes the following features:

- Infrastructure for scale-out storage.
- Design of an IBM COS solution in CD mode together with Cisco UCS C240 Rack Server
- Simplified infrastructure management with Cisco UCS Manager (UCSM)

The configuration uses the following architecture for the deployment:

- 3 x Cisco UCS C240 Rack Servers
- 2 x Cisco UCS 6332 Fabric Interconnect
- 1 x Cisco UCS Manager
- 2 x Cisco Nexus C9332PQ Switches

This solution has various options to scale capacity. The tested configuration uses an IDA (Information Dispersal Algorithm, s.) of 18/9/11 with 3 Slicestors. A base capacity summary for the tested solution is listed in Table 1 .

Table 1 Usable Capacity Options for Cisco Validated Design

HDD Type	Number of Disks	Concentrated Dispersal Mode
4 TB 7200-rpm LFF SAS drives	36	72 TB
6 TB 7200-rpm LFF SAS drives	36	108 TB
8 TB 7200-rpm LFF SAS drives	36	144 TB
10 TB 7200-rpm LFF SAS drives*	36	180 TB
12 TB 7200-rpm LFF SAS drives	36	216 TB

* Validated configuration

Technology Overview

Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) is a state-of-the-art data center platform that unites computing, network, storage access, and virtualization into a single cohesive system.

The main components of Cisco Unified Computing System are:

- **Computing** – The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers based on Intel Xeon Scalable processors. Cisco UCS servers offer the patented Cisco Extended Memory Technology to support applications with large datasets and allow more virtual machines (VM) per server.
- **Network** – The system is integrated onto a low-latency, lossless, 10/25/40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **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 virtualized environments to better support changing business and IT requirements.
- **Storage access** – The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access, the Cisco Unified Computing System can access storage over Ethernet (NFS or iSCSI), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership (TCO) and increased business agility
- Increased IT staff productivity through just-in-time provisioning and mobility support
- A cohesive, integrated system, which unifies the technology in the data center
- Industry standards supported by a partner ecosystem of industry leaders

Cisco UCS C240 Rack Server

The Cisco UCS C240 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 UCS managed environment to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

Figure 1 Cisco UCS C240 Rack Server

In response to ever-increasing computing and data-intensive real-time workloads, the enterprise-class Cisco UCS C240 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. Cisco UCS C240 M5 delivers outstanding levels of storage expandability with exceptional performance, comprised of 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-, 25- 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

The Cisco UCS C240 Rack Server is well suited for a wide range of enterprise workloads, including:

- Object Storage
- Big Data and analytics
- Collaboration
- Small and medium-sized business databases
- Virtualization and consolidation
- Storage servers
- High-performance appliances

Cisco UCS C240 Rack 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 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.

The Cisco Integrated Management Controller (IMC) delivers comprehensive out-of-band server management with support for many industry standards, including:

- Redfish Version 1.01 (v1.01)
- Intelligent Platform Management Interface (IPMI) v2.0
- Simple Network Management Protocol (SNMP) v2 and v3
- Syslog
- Simple Mail Transfer Protocol (SMTP)
- Key Management Interoperability Protocol (KMIP)
- HTML5 GUI
- HTML5 virtual Keyboard, Video, and Mouse (vKVM)
- Command-Line Interface (CLI)
- XML API

Management Software Development Kits (SDKs) and DevOps integrations exist for Python, Microsoft PowerShell, Ansible, Puppet, Chef, and more. For more information about integrations, see [Cisco DevNet](#).

The Cisco UCS C240 is Cisco Intersight ready. Cisco Intersight is a new cloud-based management platform that uses analytics to deliver proactive automation and support. By combining intelligence with automated actions, you can reduce costs dramatically and resolve issues more quickly.

Cisco UCS Virtual Interface Card 1387

The Cisco UCS Virtual Interface Card 1387 offers dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP+) 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) in a modular-LAN-on-motherboard (mLOM) form factor. The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot providing greater I/O expandability.

Figure 2 Cisco UCS Virtual Interface Card 1387



The Cisco UCS VIC 1387 provides high network performance and low latency for the most demanding applications, including:

- Big data, high-performance computing (HPC), and high-performance trading (HPT)
- Large-scale virtual machine deployments
- High-bandwidth storage targets and archives

The card is designed for the M5 generation of Cisco UCS C-Series Rack Servers and Cisco UCS S3260 dense storage servers. It includes Cisco's next-generation converged network adapter technology and offers a comprehensive feature set, so you gain investment protection for future feature software releases.

The card can present more than 256 PCIe standards-compliant interfaces to its host. These can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs).

This engine provides support for advanced data center requirements, including stateless network offloads for the following:

- Network Virtualization Using Generic Routing Encapsulation (NVGRE)
- Virtual extensible LAN (VXLAN)
- Remote direct memory access (RDMA)

The engine also offers support for performance optimization applications such as:

- Server Message Block (SMB) Direct
- Virtual Machine Queue (VMQ)
- Data Plane Development Kit (DPDK)
- Cisco NetFlow

Cisco UCS 6300 Series 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 (Figure 3). 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.

Figure 3 Cisco UCS 6300 Series Fabric Interconnect



The Cisco UCS 6300 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers, 5100 Series Blade Server Chassis, Cisco UCS C-Series Rack Servers, and Cisco UCS S-Series Storage Dense Server 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.

The Cisco UCS 6332 32-Port Fabric Interconnect is a 1-rack-unit (1RU) Gigabit Ethernet, and FCoE switch offering up to 2.56 Tbps throughput and up to 32 ports. The switch has 32 fixed 40-Gbps Ethernet and FCoE ports.

Both the Cisco UCS 6332UP 32-Port Fabric Interconnect and the Cisco UCS 6332 16-UP 40-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. The breakout feature can be configured on ports 1 to 12 and ports 15 to 26 on the Cisco UCS 6332UP fabric interconnect. Ports 17 to 34 on the Cisco UCS 6332 16-UP fabric interconnect support the breakout feature.

Cisco Nexus 9332PQ Switch

The Cisco Nexus® 9000 Series Switches (Figure 4) include both modular and fixed-port switches that are designed to overcome these challenges with a flexible, agile, low-cost, application-centric infrastructure.

Figure 4 Cisco Nexus 9332PQ Switch



The Cisco Nexus 9300 platform consists of fixed-port switches designed for top-of-rack (ToR) and middle-of-row (MoR) deployment in data centers that support enterprise applications, service provider hosting, and cloud computing environments. They are Layer 2 and 3 nonblocking 10 and 40 Gigabit Ethernet switches with up to 2.56 terabits per second (Tbps) of internal bandwidth.

The Cisco Nexus 9332PQ Switch is a 1-rack-unit (1RU) switch that supports 2.56 Tbps of bandwidth and over 720 million packets per second (mpps) across thirty-two 40-Gbps Enhanced QSFP+ ports.

All Cisco Nexus 9300 platform switches use dual-core 2.5-GHz x86 CPUs with 64-GB solid-state disk (SSD) drives and 16 GB of memory for enhanced network performance.

With the Cisco Nexus 9000 Series, organizations can quickly and easily upgrade existing data centers to carry 40 Gigabit Ethernet to the aggregation layer or to the spine (in a leaf-and-spine configuration) through advanced and cost-effective optics that enable the use of existing 10 Gigabit Ethernet fiber (a pair of multimode fiber strands).

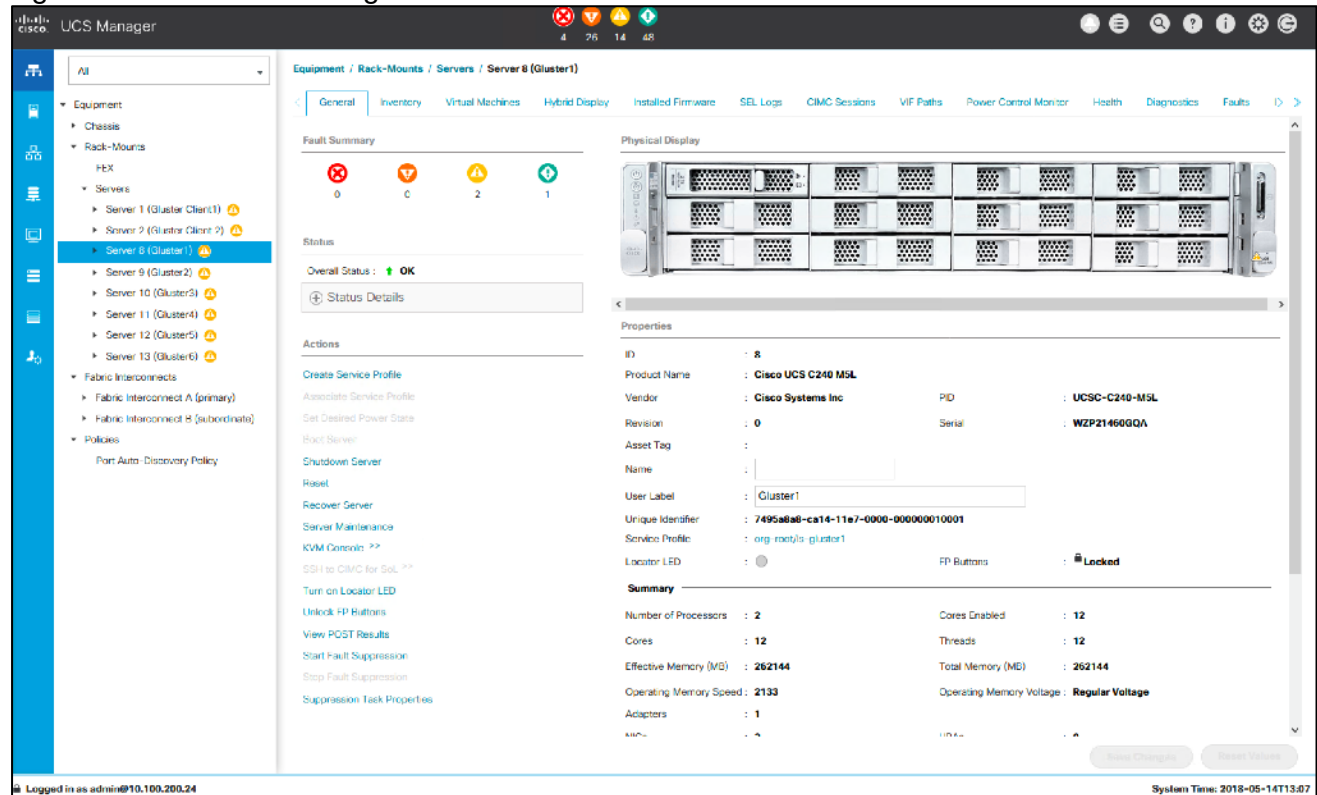
Cisco provides two modes of operation for the Cisco Nexus 9000 Series. Organizations can use Cisco NX-OS Software to deploy the Cisco Nexus 9000 Series in standard Cisco Nexus switch environments. Organizations also can use a hardware infrastructure that is ready to support Cisco Application Centric Infrastructure (Cisco ACI) to take full advantage of an automated, policy-based, systems management approach.

Cisco UCS Manager

Cisco UCS Manager (UCSM) (Figure 5) provides unified, embedded management of all software and hardware components of Cisco Unified Computing System across multiple chassis, rack servers, and thousands of virtual machines. It supports all Cisco UCS product models, including Cisco UCS B-Series Blade Servers, Cisco UCS C-

Series Rack Servers, and Cisco UCS S- and M-Series composable infrastructure and Cisco UCS Mini, as well as the associated storage resources and networks. Cisco UCS Manager is embedded on a pair of Cisco UCS 6300, 6400 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.

Figure 5 Cisco UCS Manager



An instance of Cisco UCS Manager with all Cisco UCS components, managed by Cisco UCS Manager, 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 streamlining 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. Cisco UCS 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 XML API provides comprehensive access to all Cisco UCS Manager functions. The 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 XML 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 within Cisco UCS Manager.

IBM Cloud Object Storage

The IBM COS System platform is ideal whenever enterprises need to securely store large volumes of unstructured data with high availability and where latency is not a primary consideration.

With the unprecedented growth in new digital information, use cases have emerged that enable organizations to store and distribute limitless data. A distributed and decentralized storage architecture along with an Object Storage interface enables enterprises to deliver data to their users across the globe as never before. The use cases covered in this Cisco Validated Design include:

- Content Repository
- Storage-as-a-service
- Enterprise Collaboration
- Backup
- Archive

The IBM COS System software platform uses an approach for cost-effectively storing large volumes of unstructured data while still ensuring security, availability, and reliability.

The IBM COS System storage technology uses Information Dispersal Algorithms (IDA) to separate data into unrecognizable “slices” that are distributed via network connections to storage nodes locally or across the world. The collection of dispersed storage appliances creates what is called a dispersed storage network. With dispersed storage technology, transmission and storage of data are inherently private and secure. No complete copy of the data exists in any single storage node. Only a subset of nodes needs to be available to fully retrieve the data on the network.

IDA technology transforms data into slices by using equations such that a subset of the slices can be used to re-create the original data. These slices, which are like packets but are for data storage, are then stored across multiple storage appliances (or storage nodes). Slices are created by using a combination of erasure coding, encryption, and sophisticated dispersal algorithms.

Dispersed storage systems are well-suited for storing unstructured data like digital media of all types, documents that are produced by desktop productivity applications, and server log files, which are typically larger files. Dispersal is not optimized for transaction-oriented primary storage for databases and similar high IOP workloads because of the extra processing associated with slicing and dispersing.

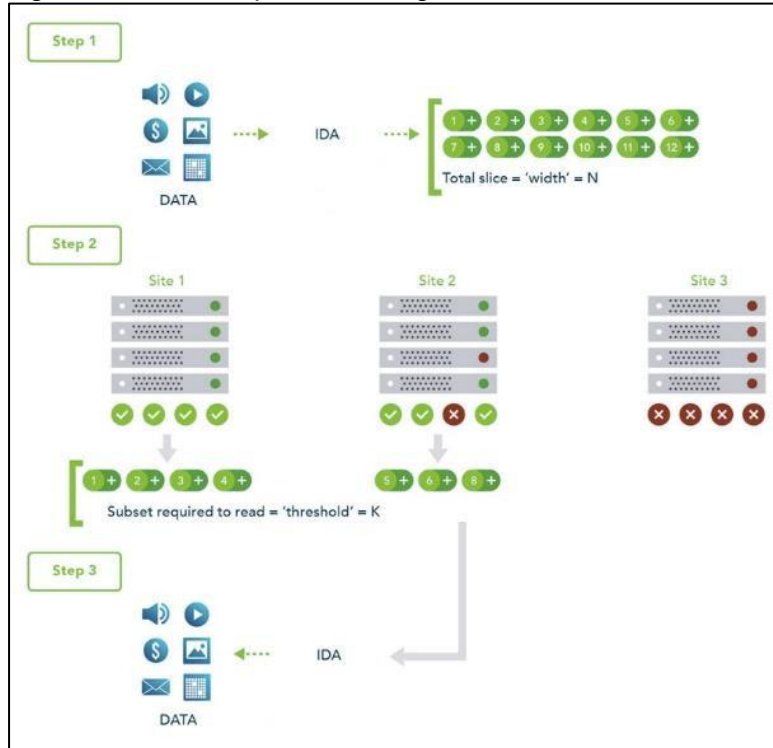
At a basic level, the IBM COS System platform uses three steps for slicing, dispersing, and retrieving data (Figure 6):

1. Data is virtualized, transformed, sliced, and dispersed by using IDAs. In the first figure, the data is separated into 12 slices. So the "width" (n) of the system is 12.
2. Slices are distributed to some combination of separate disks, storage nodes, and geographic locations. In this example, the slices are distributed to three different sites.
3. The data is retrieved from a subset of slices. In this example, the number of slices that are needed to retrieve the data is 7. So the "threshold" (k) of the system is 7.

Given a width of 12 and a threshold of 7, you can refer to this example as a "7 of 12" (k of n) configuration.

The configuration of a system is determined by the level of reliability needed. In a "7 of 12" configuration, five slices can be lost or unavailable and the data can still be retrieved because the threshold of seven slices is met. With a "5 of 8" configuration, only three slices can be lost, so the level of reliability is lower. Conversely, with a "20 of 32" configuration, 12 slices can be lost, so the level of reliability is higher.

Figure 6 How Dispersed Storage Works



Concentrated Dispersal Mode

With the Concentrated Dispersal feature used in this CVD, each SliceStor device may be responsible for multiple slices of a given object stored in the system. Previously, each SliceStor device could only hold one slice of a given object. Prior to this the IDA Width set a lower bound on the number of SliceStor devices necessary in the deployment. For example, when using an IDA Width of 18, one would have to deploy at least 18 SliceStor devices. With the Concentrated Dispersal feature, each SliceStor device might hold 6 slices of each object, and so only 3 SliceStor devices would be necessary. By continuing to use a large IDA Width, high reliability and storage efficiency can be attained, but now with a much smaller hardware footprint.

Deploying a system requires the selection of an IDA Width, Threshold, and Write Threshold that provides an acceptable level of storage efficiency, reliability, and availability. Achieving good reliability requires a sufficient difference between Threshold and Write Threshold. A high write availability, similarly requires a sufficient difference between Width and Write Threshold. Finally, for the system to provide good storage efficiency requires the ratio between IDA Width and Threshold to be as small as possible while meeting the previous reliability and availability goals.

Previous design considerations used one pillar per SliceStor and had a few impacts:

- Initial deployments were expensive which precludes some customers with small data storage requirements from using dispersed storage.
- It set a high lower-bound on the amount of usable storage capacity the system provides and expands by. This made it difficult to support sub-PB systems efficiently.

- It forces some customers with small numbers of Slicestor devices to compromise with either less reliable, or less efficient IDA configurations.

Concentrated Dispersal enables systems on the order of a few hundred terabytes, while retaining the superb reliability and storage efficiency properties that previously could only be achieved in much larger systems comprised of dozens of devices.

When the number of Slicestor appliances is an integer divisor of the IDA Width all 18 pillars still exist, but each Slicestor appliance is now responsible for more than a single pillar's worth of data. Slicestor appliances maintain an ordered namespace assignment across all the disks, which together with the safety margin, prevents the case where one disk holds more than one slice of a given object. This can only happen if a sufficient number (half or more) of drives fail within a Slicestor appliance without replacement.

Figure 7 Concentrated Dispersal Stripe-Pillar Relationship for IDA 18 used in this Design CVD



Use Cases

Concentrated Dispersal is ideally suited for the case where the system size is relatively small and is expected to remain small for the near term. There is no setting to enable or disable this feature. Instead it is activated whenever a Device Set is created which has 3 - 7 Slicestor devices. This feature should be used when the required usable capacity is on the order of 20 hard drives worth of usable storage capacity. Common reasons for deploying this quantity of Slicestor devices include:

- When the planned path of expanding the system needs to remain small (several Slicestor devices at a time)
- When seeking to minimize cost with an entry-level system

The following are some reasons to reconsider using the Concentrated Dispersal feature:

- When it is necessary to use a feature that Concentrated Dispersal is not compatible with like Container mode or partially populated Slicestors
- When non-standard IDA Configurations must be explicitly configured like less pillars than drives
- When maintenance operations cannot readily replace failed hard drives
- When high performance is a primary concern
- When achieving the highest possible availability is a paramount concern

Cloud Object Storage Components

You can use the IBM COS System platform to create storage systems that have three software components: the IBM COS Manager software, IBM COS Accesser software and IBM COS Slicestor software.

The software components can be deployed on a wide range of compatible industry-standard hardware platforms, as virtual machines, and in the case of the IBM COS Accesser software, as a software application that is running on a Linux operating system. Physical and virtual deployment can be combined in a single system by using virtual machine deployment for the IBM COS Manager and IBM COS Accesser software and physical servers for the IBM COS Slicestor software as an example.

Each of the three software components serves a specific function:

- The IBM COS Manager software is responsible for monitoring the health and performance of the system, configuring the system and provisioning storage, managing faults, and other administrative and operational functions.
- The IBM COS Accesser software is responsible for encrypting/encoding data on ingest and decoding/decrypting it when read and managing the dispersal of slices of data resulting from this process across a set of IBM COS Slicestor nodes.
- The IBM COS Slicestor software is responsible for the storage of slices.

The underlying storage pool of a dispersed or concentrated dispersed storage system can be shared and is jointly accessible by multiple access protocols.

Object-based Access Methods

The Simple Object interface is accessed with a HTTP/REST API. Simple PUT, GET, DELETE, and LIST commands allow applications to access digital content, and the resulting object ID is stored directly within the application. With the IBM COS Accesser application, no IBM COS Accesser appliance is needed since the application server can talk directly to IBM COS Slicestor storage nodes.

REST API Access to Storage

Figure 8 REST API Storage Interfaces



REST is a style of software architecture for distributed hypermedia information retrieval systems such as the World Wide Web. REST-style architectures consist of clients and servers. Clients initiate requests to servers. Servers process requests and return associated responses. Requests and responses are built around the transfer of various representations of the resources.

The REST API works in way that is similar to retrieving a Universal Resource Locator (URL). But instead of requesting a webpage, the application is referencing an object.

REST API access to storage offers several advantages:

- Tolerates internet latency
- Provides for "programmable" storage
- Provides efficient global access to large amounts of data

Embedded Accesser

This CVD uses an Embedded Accesser Appliance function. The Embedded Accesser Appliance feature provides Accesser Appliance functions on the IBM COS Slicestor Appliance. This feature provides customers an opportunity to save on capital expenses by using one physical appliance for both Accesser and Slicestor appliance functions. However, before you deploy this feature, careful consideration needs to be given to the Slicestor hardware and the workload presented to the servers and the load balancing between the available Slicestor appliances.

Network

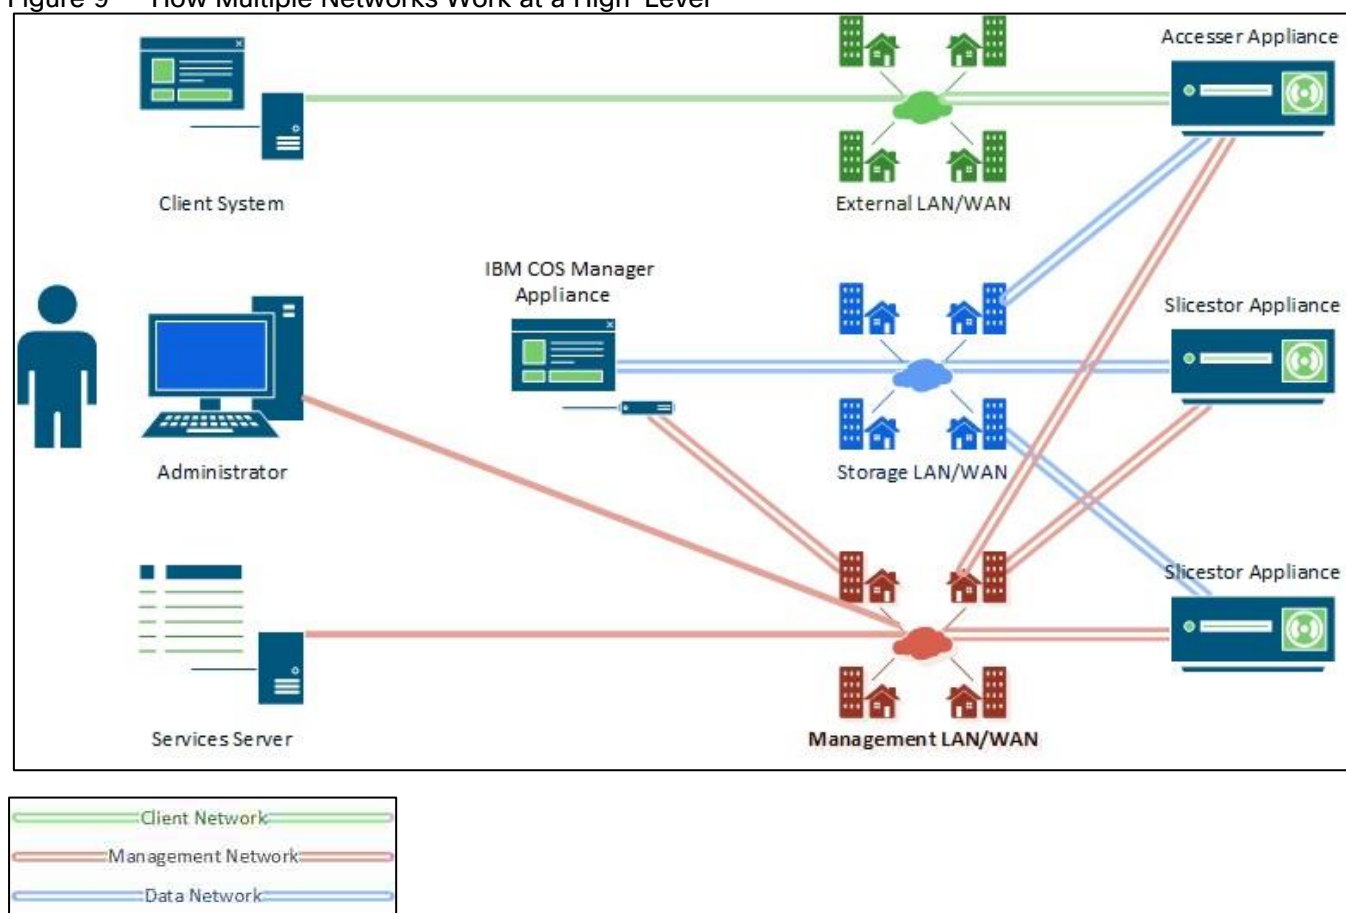
Network administrators can configure the first four layers of the OSI model on a system to use separate network traffic between storage data, management information, and client data.

An IBM COS System that uses certified devices, can dedicate network interfaces (NICs) to three distinct networks to transfer:

- Data within the system
- Management information to management systems
- Data to a client application

These networks are referred to as channels.

Figure 9 How Multiple Networks Work at a High-Level



In separating data into channels, the system provides better security, more flexible management options and minimizes network congestion for high-performance applications.

IBM Hardware Requirements

Hardware with a minimum of 10 GbE interconnect and a RAM capacity of 256 GB is recommended for a full-scale deployment of an IBM COS System with Embedded Accesser Appliance functions.

For Concentrated Dispersal Mode the following requirements are important:

- Slicestor requirements
 - Slicestor appliances in a Concentrated Dispersal system must have at least 12 Drives
 - Slicestor appliances in a Concentrated Dispersal system must have at least 32 GB of memory
 - Slicestor appliances in a Concentrated Dispersal system using embedded Accessers must have at least 128 GB of memory
- Device set requirements
 - Device sets must have between 3 and 6 Slicestor devices
 - System expansion requires the new set to equal the **Concentrated Dispersal** set size or an integer multiple of the full IDA width
- Device evacuation requirements
 - Device evacuation requires destination Slicestor appliances to have an equal or greater number of drives
- Mirroring requirements
 - Vault mirrors must use the same vault optimization on both sides of the mirror
 - Vault mirrors must use Concentrated Dispersal Device Set sizes (3 - 6 devices) on both sides of the mirror

Solution Design

Solution Overview


This Cisco Validated Design provides a comprehensive, end-to-end guide for deploying IBM COS in Concentrated Dispersal Mode with Embedded Accesser on Cisco UCS C240 within infrastructure made possible by Cisco UCS Manager and the Cisco UCS 6332 Fabric Interconnects.

One of the key design goals of this scale out architecture was to deploy all elements on 40GbE networking end to end within a single Cisco UCS domain and start small with CD mode for IBM COS. Both IBM COS components – Embedded Accesser, and SliceStor – utilize the robust throughput and low latency only provided by the Cisco UCS 6332 Fabric Interconnect. Additionally, both components take advantage of the flexibility provided by the stateless nature of Cisco UCS service profiles and service profile templates.

This design uses the Cisco Nexus 9000 series data center switches in NX-OS standalone mode but provides investment protection to migrate to ACI or higher network bandwidths (1/10/25/40/50/100Gbps) while enabling innovative analytics and visibility using Tetration and automation that support in-box and off-box Python scripting and Open NX-OS that support dev-ops tools (Chef, Puppet, Ansible).

The key design for IBM COS with CD mode on Cisco UCS C240 is shown in Figure 10 .

Figure 10 Topology of IBM COS with CD Mode on Cisco UCS C240 M5L



- Manager instance deployed as virtual machine OVA
- Embedded Accesser deployed on SliceStor
- SliceStor deployed on Cisco UCS C240
- Cisco UCS C240 connected to UCS 6332 Fabric Interconnect with 40Gbps line speed
- Cisco UCS 6332 Fabric Interconnect connected to Nexus 9332PQ with 40Gbps line speed

IBM Cloud Object Configuration for Cisco Validated Design

The current Design Guide uses the following configuration for IBM COS:

- Virtual Manager using an OVA
- Embedded Accesser
- Concentrated Dispersal Mode with IDA 18/9/11

Details for the specific CD Mode are listed in Table 2 .

Table 2 Configuration for CD Mode used in this Design Guide

Sites	IDA	Store Count	Disks per Store	Usable Capacity	Expansion Factor	Availability	Reliability
1	18/9/11	3	36	180 TB	2.00	5 nines	15+ nines

General Hardware Requirements

Table 3 List of Components

Component	Model	Quantity	Comments
IBM Slicestor/Accesser	Cisco UCS C240	3	Per server node: <ul style="list-style-type: none"> • 2 x Intel Skylake 4xxx/5xxx • 128-256 GB Memory • 1 x VIC 1380 • 12 Gbit SAS RAID Controller • Disks <ul style="list-style-type: none"> ◦ 2 x SSD/HDD RAID 1 – Boot ◦ 12 x NL-SAS HDD JBOD – Data
IBM Manager	Virtual Machine OVA	1	<ul style="list-style-type: none"> • 4 vCPU • 16 GB Memory • 128 GB Disk • 1 x Network
Cisco UCS Fabric Interconnects	Cisco UCS 6332 Fabric Interconnects	2	
Switches	Cisco Nexus 9332PQ	2	

Compute Layer Design

Each Cisco UCS C240 rackmount server is equipped with a Cisco UCS Virtual Interface Card (VIC) supporting dual 40-Gbps fabric connectivity. The Cisco UCS VICs eliminate the need for separate physical interface cards on each server for data and management connectivity. For this solution with IBM COS ClevOS the VIC is configured with two virtual NICs, one on each physical VIC interface. IBM COS is configured to leverage these two vNICs to provide operational active-backup redundancy in software.

Cisco UCS Server Connectivity to Unified Fabric

Cisco UCS servers are typically deployed with a single VIC card for unified network and storage access. The Cisco VIC connects into a redundant unified fabric provided by a pair of Cisco UCS Fabric Interconnects. Fabric Interconnects are an integral part of Cisco Unified Computing System, providing unified management and connectivity to all attached blades, chassis and rack servers. Fabric Interconnects provide a lossless and

deterministic FCoE fabric. For the servers connected to it, the Fabric Interconnects provide LAN, SAN and management connectivity to the rest of the network.

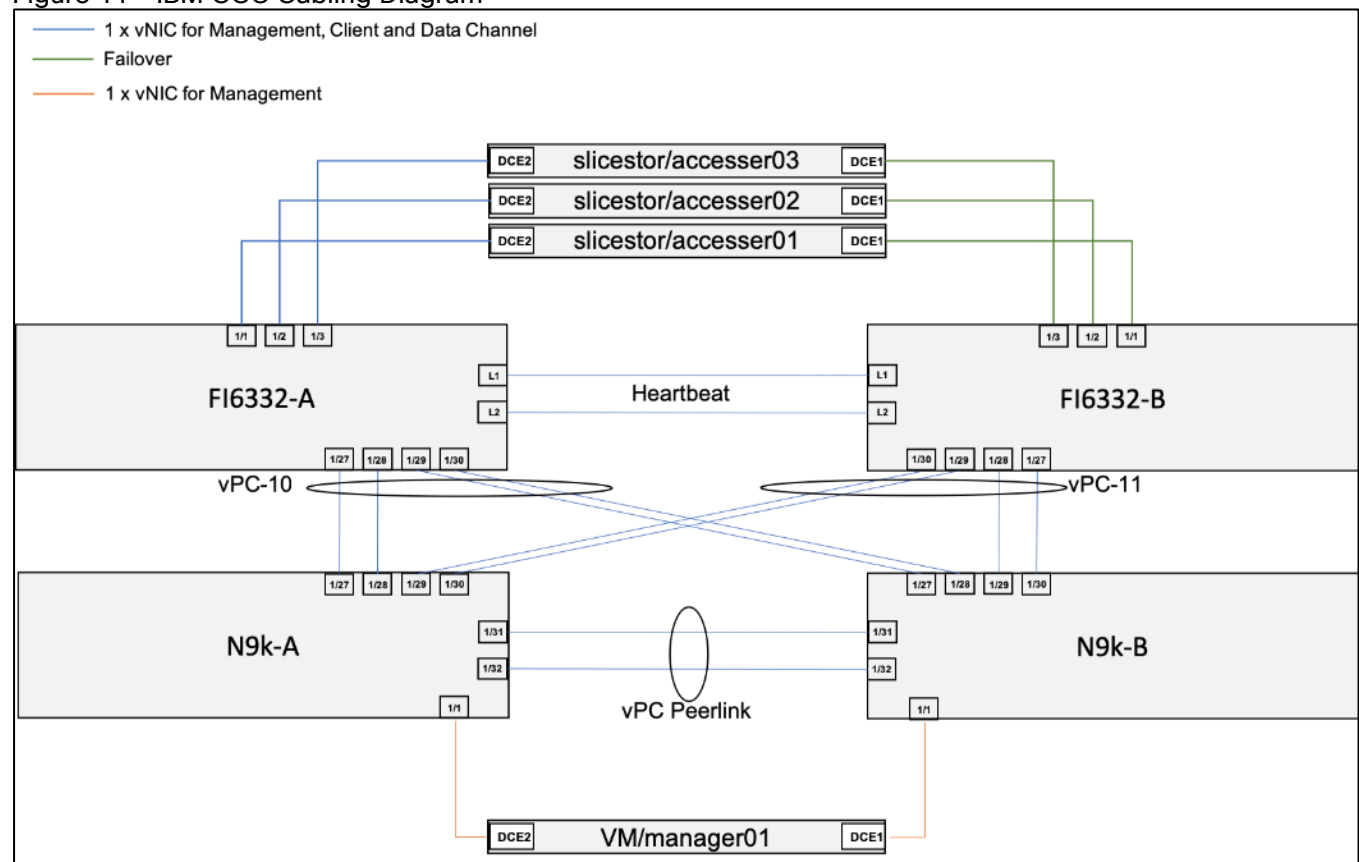
Validated Compute Design

The connectivity of the solution is based on 40 Gbps. All components are connected together via 40 Gbps QSFP cables except the virtual Manager node, which uses a 10 Gbit connectivity. Between both Cisco Nexus 9332PQ switches are 2 x 40 Gbit cabling. Each Cisco UCS 6332 Fabric Interconnect is connected via 2 x 40 Gbps to each Cisco UCS 9332PQ switch. And each Cisco UCS C240 M5L rack server is connected with a single 40 Gbit cable to each Fabric Interconnect.

The exact cabling for the IBM COS solution is illustrated in following picture. It shows also the vNIC configuration for the Data, Client and Management channel.

The virtual IBM COS Management node is connected to both Nexus N9k-A and N9k-B and has access to the Slicestors/Accessers.

Figure 11 IBM COS Cabling Diagram



For a better reading and overview, the exact physical connectivity between the Cisco UCS 6332 Fabric Interconnects and the Cisco UCS C-Class server is listed in Table 4 .

Table 4 Physical Connectivity between FI 6332 and Cisco UCS C240 M5L

Port	Role	FI6332-A	FI6332-B
Eth1/1	Server	slicestor/accesser01, DCE2	slicestor/accesser01, DCE1
Eth1/2	Server	slicestor/accesser02, DCE2	slicestor/accesser02, DCE1

Port	Role	FI6332-A	FI6332-B
Eth1/3	Server	slice stor/accesser03, DCE2	slice stor/accesser03, DCE1
Eth1/27	Network	N9k-A, Eth1/27	N9k-B, Eth1/30
Eth1/28	Network	N9k-A, Eth1/28	N9k-B, Eth1/29
Eth1/29	Network	N9k-B, Eth1/27	N9k-A, Eth1/30
Eth1/30	Network	N9k-B, Eth1/28	N9k-B, Eth1/29

High Availability

The Cisco and IBM solution was designed for maximum availability of the complete infrastructure (compute, network, storage) with no single points of failure.

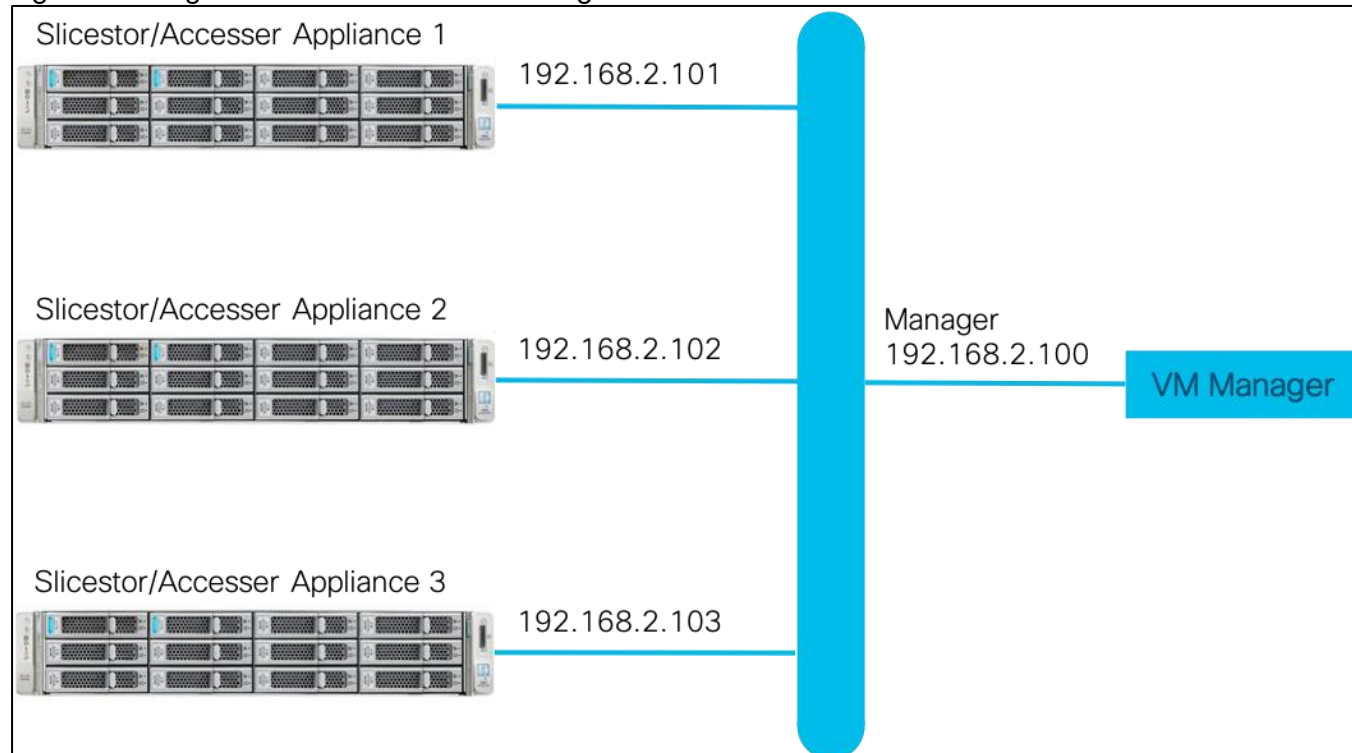
Compute

- Cisco UCS provides redundancy at the component and link level and end-to-end path redundancy to the LAN network.
- Cisco UCS C240 Rack Server is highly redundant with redundant power supplies and fans.
- Each server is deployed using vNICs that provide redundant connectivity to the unified fabric. NIC failover is enabled between Cisco UCS Fabric Interconnects using Cisco UCS Manager. This is done for all Slice stor with Embedded Accesser node vNICs.

Network

- Link aggregation using port channels and virtual port channels can be used throughout the design for higher bandwidth and availability, if the optional Cisco UCS Nexus 9332 is deployed. Between each Cisco UCS 6332 Fabric Interconnect and both Cisco Nexus 9332PQ is one virtual Port Channel (vPC) configured. vPCs allow links that are physically connected to two different Cisco Nexus 9000 switches to appear to the Fabric Interconnect as coming from a single device and as part of a single port channel.
- Each Slice stor with Embedded Accesser is configured in mode 1 active-backup bonding mode at the ClevOS software layer.

Figure 12 illustrates the logical configuration of the network for the IBM COS solution with Embedded Accesser and CD mode.

Figure 12 Logical View of the Network Configuration used in this CVD

QoS and Jumbo Frames

Cisco UCS, Cisco Nexus, and IBM COS nodes in this solution provide QoS policies and features for handling congestion and traffic spikes. The network-based QoS capabilities in these components can alleviate and provide the priority that the different traffic types require.

This design also recommends end-to-end jumbo frames with an MTU of 9000 Bytes across the LAN and Unified Fabric links. Jumbo frames increase the throughput between devices by enabling larger sized frames to be sent and received on the wire while reducing the CPU resources necessary to process them. Jumbo frames were enabled during validation on the LAN network links in the Cisco Nexus switching layer and on the Unified Fabric links.

Software Distributions and Versions

The required software distribution versions are listed below in Table 5 .

Table 5 Software Versions

Layer	Component	Version or Release
Cisco UCS C240	Adapter	4.3(1b)
	BIOS	C240M5.4.0.1c
	Board Controller	38.0
	CIMC Controller	4.0(1a)
	Storage Controller SAS 2	50.1.0-1456

Layer	Component	Version or Release
Network 6332 Fabric Interconnect	UCS Manager	4.0(1a)
	Kernel	5.0(3)N2(4.01a)
	System	5.0(3)N2(4.01a)
Network Nexus 9332PQ	BIOS	07.59
	NXOS	7.0(3)I5(1)
Software	IBM COS	3.3.16

Deployment Hardware and Software

Fabric Configuration

This section provides the details to configure a fully redundant, highly available Cisco UCS 6332 fabric configuration.

- Initial setup of Cisco Nexus C9332PQ Switch A and B
- Initial setup of the Cisco UCS Fabric Interconnect 6332 A and B
- Connect to Cisco UCS Manager using virtual IP address of the web browser
- Launch Cisco UCS Manager
- Enable server and uplink ports
- Start discovery process
- Create pools and policies for service profile template
- Create storage profiles
- Create Service Profile templates and appropriate Service Profiles
- Associate Service Profiles to servers

Configure Cisco Nexus C9332PQ Switch A and B

Both Cisco UCS Fabric Interconnect A and B are connected to two Cisco Nexus C9332PQ switches for connectivity to applications and clients. The following sections describe the setup of both Cisco Nexus C9332PQ switches.

Initial Setup of Cisco Nexus C9332PQ Switch A and B

To configure Switch A, connect a Console to the Console port of each switch, power on the switch and follow these steps:

1. Type **yes**.
2. Type **n**.
3. Type **n**.
4. Type **n**.
5. Enter the switch name.
6. Type **y**.
7. Type your IPv4 management address for Switch A.
8. Type your IPv4 management netmask for Switch A.

9. Type `y`.
10. Type your IPv4 management default gateway address for Switch A.
11. Type `n`.
12. Type `n`.
13. Type `y` for ssh service.
14. Press `<Return>` and then `<Return>`.
15. Type `y` for ntp server.
16. Type the IPv4 address of the NTP server.
17. Type `in L2` for interface layer.
18. Press `<Return>` and again `<Return>`.
19. Check the configuration and if correct then press `<Return>` and again `<Return>`.

The complete setup looks like the following:

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

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

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

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

```
----- 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.

Please register Cisco Nexus9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus9000 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 : N9k-A
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
    Mgmt0 IPv4 address : 10.100.200.20
    Mgmt0 IPv4 netmask : 255.255.255.0
  Configure the default gateway? (yes/no) [y]:
    IPv4 address of the default gateway : 10.100.200.1
  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> [1024]:
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : 10.100.200.31
  Configure default interface layer (L3/L2) [L3]: L2
  Configure default switchport interface state (shut/noshut) [shut]:
  Configure CoPP system profile (strict/moderate/lenient/dense) [strict]:
The following configuration will be applied:
  password strength-check
  switchname N9k-A
vrf context management
ip route 0.0.0.0/0 10.100.200.1
exit
  no feature telnet
  ssh key rsa 1024 force
  feature ssh
```

```

ntp server 10.100.200.31

no system default switchport

system default switchport shutdown

copp profile strict

interface mgmt0

ip address 10.100.200.20 255.255.255.0

no shutdown

```

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

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

[#####] 100%

Copy complete.

User Access Verification

N9k-A login:



Repeat the same steps for the Cisco Nexus 9332PQ Switch B with the exception of configuring a different IPv4 management address in step 7.

Enable Features on Cisco Nexus 9332PQ Switch A and B

To enable the features UDLD, VLAN, LACP, VPC, and Jumbo Frames, connect to the management interface via ssh on both switches and follow these steps on both Switch A and B:

Switch A

```

N9k-A# configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-A(config)# feature udld

N9k-A(config)# feature interface-vlan

N9k-A(config)# feature lacp

N9k-A(config)# feature vpc

N9k-A(config)# system jumbomtu 9216

N9k-A(config)# spanning-tree port type edge bpduguard default

N9k-A(config)# spanning-tree port type edge bpdufilter default

```

```

N9k-A(config)# port-channel load-balance src-dst ip-l4port-vlan
N9k-A(config)# exit
N9k-A#

```

Switch B

```

N9k-B# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
N9k-B(config)# feature udld
N9k-B(config)# feature interface-vlan
N9k-B(config)# feature lacp
N9k-B(config)# feature vpc
N9k-B(config)# system jumbomtu 9216
N9k-B(config)# spanning-tree port type edge bpduguard default
N9k-B(config)# spanning-tree port type edge bpdufilter default
N9k-B(config)# port-channel load-balance src-dst ip-l4port-vlan
N9k-B(config)# exit
N9k-B#

```

Configuring VLANs on Nexus 9332PQ Switch A and B

To configure VLAN Native-VLAN and Public-VLAN, follow these steps on Switch A and Switch B:

Switch A

```

N9k-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
N9k-A(config)# vlan 2
N9k-A(config-vlan)# name Native-VLAN
N9k-A(config-vlan)# exit
N9k-A(config)# vlan 10
N9k-A(config-vlan)# name Public-VLAN
N9k-A(config-vlan)# exit

```

Switch B

```

N9k-B# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
N9k-B(config)# vlan 2
N9k-B(config-vlan)# name Native-VLAN

```



```

N9k-B(config-vlan)# exit
N9k-B(config)# vlan 10
N9k-B(config-vlan)# name Public-VLAN
N9k-B(config-vlan)# exit

```

Configuring vPC Domain on Nexus 9332PQ Switch A and B

To configure the vPC Domain, follow these steps on Switch A and Switch B:

Switch A

```

N9k-A# config terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-A(config)# vpc domain 1

N9k-A(config-vpc-domain)# role priority 10

N9k-A(config-vpc-domain)# peer-keepalive destination 10.100.200.21 source
10.100.200.20

N9k-A(config-vpc-domain)# peer-switch

N9k-A(config-vpc-domain)# peer-gateway

N9k-A(config-vpc-domain)# ip arp synchronize

N9k-A(config-vpc-domain)# auto-recovery

N9k-A(config-vpc-domain)# copy run start

N9k-A(config-vpc-domain)# exit

```

Switch B

```

N9k-B# config terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-B(config)# vpc domain 1

N9k-B(config-vpc-domain)# role priority 20

N9k-B(config-vpc-domain)# peer-keepalive destination 10.100.200.20 source
10.100.200.21

N9k-B(config-vpc-domain)# peer-switch

N9k-B(config-vpc-domain)# peer-gateway

N9k-B(config-vpc-domain)# ip arp synchronize

N9k-B(config-vpc-domain)# auto-recovery

N9k-B(config-vpc-domain)# copy run start

N9k-B(config-vpc-domain)# exit

```

Configuring Network Interfaces for vPC Peer Links on Nexus 9332PQ Switch A and B

To configure the network interfaces for vPC Peer Links, follow these steps on Switch A and Switch B:

Switch A

```
N9k-A# config terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-A(config)# interface Eth 1/31
N9k-A(config-if)# description VPC Peer Nexus B Port 1/31
N9k-A(config-if)# interface Eth 1/32
N9k-A(config-if)# description VPC Peer Nexus B Port 1/32
N9k-A(config-if)# interface Eth1/31,Eth1/32
N9k-A(config-if)# channel-group 1 mode active
N9k-A(config-if)# no shutdown
N9k-A(config-if)# udld enable
N9k-A(config-if)# interface port-channel 1
N9k-A(config-if)# description vPC peer-link
N9k-A(config-if)# switchport
N9k-A(config-if)# switchport mode trunk
N9k-A(config-if)# switchport trunk native vlan 2
N9k-A(config-if)# switchport trunk allowed vlan 10
N9k-A(config-if)# spanning-tree port type network
N9k-A(config-if)# vpc peer-link
N9k-A(config-if)# no shutdown
N9k-A(config-if)# copy run start
```

Switch B

```
N9k-A# config terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-B(config)# interface Eth 1/31
N9k-B(config-if)# description VPC Peer Nexus A Port 1/31
N9k-B(config-if)# interface Eth 1/32
N9k-B(config-if)# description VPC Peer Nexus A Port 1/32
N9k-B(config-if)# interface Eth1/31,Eth1/32
N9k-B(config-if)# channel-group 1 mode active
```

```

N9k-B(config-if)# no shutdown
N9k-B(config-if)# uddl enable
N9k-B(config-if)# interface port-channel 1
N9k-B(config-if)# description vPC peer-link
N9k-B(config-if)# switchport
N9k-B(config-if)# switchport mode trunk
N9k-B(config-if)# switchport trunk native vlan 2
N9k-B(config-if)# switchport trunk allowed vlan 10
N9k-B(config-if)# spanning-tree port type network
N9k-B(config-if)# vpc peer-link
N9k-B(config-if)# no shutdown
N9k-B(config-if)# copy run start

```

Configuring Network Interfaces to Cisco UCS FI 6332 on Nexus 9332PQ Switch A and B

To configure the network interfaces to Cisco UCS FI 6332, follow these steps on Switch A and Switch B:

Switch A

```

N9k-A(config-if)# interface port-channel 10
N9k-A(config-if)# description Port Channel FI-A
N9k-A(config-if)# switchport
N9k-A(config-if)# switchport mode trunk
N9k-A(config-if)# switchport trunk native vlan 2
N9k-A(config-if)# switchport trunk allowed vlan 10
N9k-A(config-if)# spanning-tree port type edge trunk
N9k-A(config-if)# spanning-tree guard root
N9k-A(config-if)# no lacp graceful-convergence
N9k-A(config-if)# mtu 9216
N9k-A(config-if)# vpc 10
N9k-A(config-if)# no shutdown
N9k-A(config-if)# interface Eth1/27, Eth 1/28
N9k-A(config-if)# description Interface Port Channel FI-A
N9k-A(config-if)# switchport
N9k-A(config-if)# switchport mode trunk
N9k-A(config-if)# switchport trunk native vlan 2

```

```
N9k-A(config-if)# switchport trunk allowed vlan 10
N9k-A(config-if)# mtu 9216
N9k-A(config-if)# channel-group 10 mode active
N9k-A(config-if)# no shutdown
N9k-A(config-if)# udld enable

N9k-A(config-if)# interface port-channel 11
N9k-A(config-if)# description Port Channel FI-B
N9k-A(config-if)# switchport
N9k-A(config-if)# switchport mode trunk
N9k-A(config-if)# switchport trunk native vlan 2
N9k-A(config-if)# switchport trunk allowed vlan 10
N9k-A(config-if)# spanning-tree port type edge trunk
N9k-A(config-if)# spanning-tree guard root
N9k-A(config-if)# no lacp graceful-convergence
N9k-A(config-if)# mtu 9216
N9k-A(config-if)# vpc 11
N9k-A(config-if)# no shutdown
N9k-A(config-if)# interface Eth1/29, Eth 1/30
N9k-A(config-if)# description Interface Port Channel FI-B
N9k-A(config-if)# switchport
N9k-A(config-if)# switchport mode trunk
N9k-A(config-if)# switchport trunk native vlan 2
N9k-A(config-if)# switchport trunk allowed vlan 10
N9k-A(config-if)# mtu 9216
N9k-A(config-if)# channel-group 11 mode active
N9k-A(config-if)# no shutdown
N9k-A(config-if)# udld enable
N9k-A(config-if)# copy run start
```

Switch B

```
N9k-B(config-if)# interface port-channel 10
N9k-B(config-if)# description Port Channel FI-A
```

```
N9k-B(config-if)# switchport
N9k-B(config-if)# switchport mode trunk
N9k-B(config-if)# switchport trunk native vlan 2
N9k-B(config-if)# switchport trunk allowed vlan 10
N9k-B(config-if)# spanning-tree port type edge trunk
N9k-B(config-if)# spanning-tree guard root
N9k-B(config-if)# no lacp graceful-convergence
N9k-B(config-if)# mtu 9216
N9k-B(config-if)# vpc 10
N9k-B(config-if)# no shutdown
N9k-B(config-if)# interface Eth1/27, Eth 1/28
N9k-B(config-if)# description Interface Port Channel FI-A
N9k-B(config-if)# switchport
N9k-B(config-if)# switchport mode trunk
N9k-B(config-if)# switchport trunk native vlan 2
N9k-B(config-if)# switchport trunk allowed vlan 10
N9k-B(config-if)# mtu 9216
N9k-B(config-if)# channel-group 10 mode active
N9k-B(config-if)# no shutdown
N9k-B(config-if)# udld enable

N9k-B(config-if)# interface port-channel 11
N9k-B(config-if)# description Port Channel FI-B
N9k-B(config-if)# switchport
N9k-B(config-if)# switchport mode trunk
N9k-B(config-if)# switchport trunk native vlan 2
N9k-B(config-if)# switchport trunk allowed vlan 10
N9k-B(config-if)# spanning-tree port type edge trunk
N9k-B(config-if)# spanning-tree guard root
N9k-B(config-if)# no lacp graceful-convergence
N9k-B(config-if)# mtu 9216
N9k-B(config-if)# vpc 11
```

```

N9k-B(config-if)# no shutdown
N9k-B(config-if)# interface Eth1/29, Eth 1/30
N9k-B(config-if)# description Interface Port Channel FI-B
N9k-B(config-if)# switchport
N9k-B(config-if)# switchport mode trunk
N9k-B(config-if)# switchport trunk native vlan 2
N9k-B(config-if)# switchport trunk allowed vlan 10
N9k-B(config-if)# mtu 9216
N9k-B(config-if)# channel-group 11 mode active
N9k-B(config-if)# no shutdown
N9k-B(config-if)# uddld enable
N9k-B(config-if)# copy run start

```

Verification Check of Cisco Nexus C9332PQ Configuration for Switch A and B

Switch A

```
N9k-B# config terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
N9k-A(config)# show vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```

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

```

```

Delay-restore status          : Timer is off.(timeout = 150s)
Delay-restore SVI status      : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Disabled

```

vPC Peer-link status

```

-----
id      Port      Status Active vlans
--      -
1       Po1       up      2,10

```

vPC status

```

-----
Id      Port      Status Consistency Reason          Active vlans
--      -
10      Po10      up      success      success      2,10
11      Po11      up      success      success      2,10

```

Please check "show vpc consistency-parameters vpc <vpc-num>" for the consistency reason of down vpc and for type-2 consistency reasons for any vpc.

N9k-A(config)# show port-channel summary

```

Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lACP mode (member)
        M - Not in use. Min-links not met

```

Group	Port-Channel	Type	Protocol	Member	Ports
1	Po1 (SU)	Eth	LACP	Eth1/31 (P)	Eth1/32 (P)
10	Po10 (SU)	Eth	LACP	Eth1/29 (P)	Eth1/30 (P)
11	Po11 (SU)	Eth	LACP	Eth1/27 (P)	Eth1/28 (P)

Switch B

N9k-B# config terminal

Enter configuration commands, one per line. End with CNTL/Z.

N9k-A(config)# show vpc brief

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```

vPC domain id                : 1
Peer status                   : peer adjacency formed ok
vPC keep-alive status         : peer is alive
Configuration consistency status : success
Per-vlan consistency status    : success
Type-2 consistency status      : success
vPC role                      : primary
Number of vPCs configured      : 2
Peer Gateway                   : Enabled
Dual-active excluded VLANs     : -
Graceful Consistency Check     : Enabled
Auto-recovery status          : Enabled, timer is off.(timeout = 240s)
Delay-restore status           : Timer is off.(timeout = 150s)
Delay-restore SVI status       : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Disabled

```

vPC Peer-link status

id	Port	Status	Active vlans
----	------	--------	--------------

1	Po1	up	2,10
---	-----	----	------

vPC status

Id	Port	Status	Consistency	Reason	Active vlans
----	------	--------	-------------	--------	--------------

10	Po10	up	success	success	2,10
----	------	----	---------	---------	------

11	Po11	up	success	success	2,10
----	------	----	---------	---------	------

Please check "show vpc consistency-parameters vpc <vpc-num>" for the consistency reason of down vpc and for type-2 consistency reasons for any vpc.

N9k-A(config)# show port-channel summary

Flags: D - Down P - Up in port-channel (members)

I - Individual H - Hot-standby (LACP only)

s - Suspended r - Module-removed

b - BFD Session Wait

S - Switched R - Routed

U - Up (port-channel)

p - Up in delay-lACP mode (member)

M - Not in use. Min-links not met

Group	Port-Channel	Type	Protocol	Member Ports
-------	--------------	------	----------	--------------

1	Po1 (SU)	Eth	LACP	Eth1/31 (P) Eth1/32 (P)
---	----------	-----	------	-------------------------

10	Po10 (SU)	Eth	LACP	Eth1/29 (P) Eth1/30 (P)
----	-----------	-----	------	-------------------------

```

11      Po11 (SU)      Eth      LACP      Eth1/27 (P)      Eth1/28 (P)

```

The formal setup for the Cisco Nexus C9332PQ switches is now finished. The next step is to configure the Cisco UCS Fabric Interconnect 6332.

Initial Setup of Cisco UCS 6332 Fabric Interconnects

This section describes the initial setup of the Cisco UCS 6332 Fabric Interconnects A and B

Configure Fabric Interconnect A

To configure Fabric Interconnect A, follow these steps:

1. Connect to the console port on the first Cisco UCS 6332 Fabric Interconnect.
2. At the prompt to enter the configuration method, enter `console` to continue.
3. If asked to either perform a new setup or restore from backup, enter `setup` to continue.
4. Enter `y` to continue to set up a new Fabric Interconnect.
5. Enter `n` to enforce strong passwords.
6. Enter the password for the admin user.
7. Enter the same password again to confirm the password for the admin user.
8. When asked if this fabric interconnect is part of a cluster, answer `y` to continue.
9. Enter `A` for the switch fabric.
10. Enter the cluster name `FI6332` for the system name.
11. Enter the Mgmt0 IPv4 address.
12. Enter the Mgmt0 IPv4 netmask.
13. Enter the IPv4 address of the default gateway.
14. Enter the cluster IPv4 address.
15. To configure DNS, answer `y`.
16. Enter the DNS IPv4 address.
17. Answer `y` to set up the default domain name.
18. Enter the default domain name.
19. Review the settings that were printed to the console, and if they are correct, answer `yes` to save the configuration.

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

Example Setup for Fabric Interconnect A

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

This setup utility will guide you through the basic configuration of the system. Only minimal configuration including IP connectivity to the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.

To back track or make modifications to already entered values, complete input till end of section and answer no when prompted to apply configuration.

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

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

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

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

Enter the password for "admin":

Confirm the password for "admin":

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

Enter the switch fabric (A/B): A

Enter the system name: FI6332

Physical Switch Mgmt0 IP address : 10.100.200.22

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.100.200.1

Cluster IPv4 address : 10.100.200.24

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

DNS IP address : 8.8.8.8

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

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

Following configurations will be applied:

```
Switch Fabric=A
System Name=FI6332
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.100.200.22
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.100.200.1
Ipv6 value=0
DNS Server=8.8.8.8
```

```
Cluster Enabled=yes
Cluster IP Address=10.100.200.24
```

NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.

UCSM will be functional only after peer FI is configured in clustering mode.

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

```
Applying configuration. Please wait.
```

```
Configuration file - Ok
```

```
Cisco UCS 6300 Series Fabric Interconnect
```

```
FI6332-A login:
```

Configure Fabric Interconnect B

To configure Fabric Interconnect B, follow these steps:

1. Connect to the console port on the second Cisco UCS 6332 Fabric Interconnect.
2. When prompted to enter the configuration method, enter `console` to continue.
3. The installer detects the presence of the partner Fabric Interconnect and adds this fabric interconnect to the cluster. Enter `y` to continue the installation.
4. Enter the admin password that was configured for the first Fabric Interconnect.

5. Enter the Mgmt0 IPv4 address.
6. Answer `yes` to save the configuration.
7. Wait for the login prompt to confirm that the configuration has been saved.

Example Setup for Fabric Interconnect B

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

This setup utility will guide you through the basic configuration of the system. Only minimal configuration including IP connectivity to the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.
To back track or make modifications to already entered values, complete input till end of section and answer no when prompted to apply configuration.

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

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

Enter the admin password of the peer Fabric interconnect:

Connecting to peer Fabric interconnect... done

Retrieving config from peer Fabric interconnect... done

Peer Fabric interconnect Mgmt0 IPv4 Address: 10.100.200.22

Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0

Cluster IPv4 address : 10.100.200.24

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

Physical Switch Mgmt0 IP address : 10.100.200.23

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

```
Applying configuration. Please wait.
```

```
Fri Sep 30 05:41:48 UTC 2016
```

```
Configuration file - Ok
```

```
Cisco UCS 6300 Series Fabric Interconnect
```

```
FI6332-B login:
```

Logging Into Cisco UCS Manager

To login to Cisco UCS Manager, follow these steps:

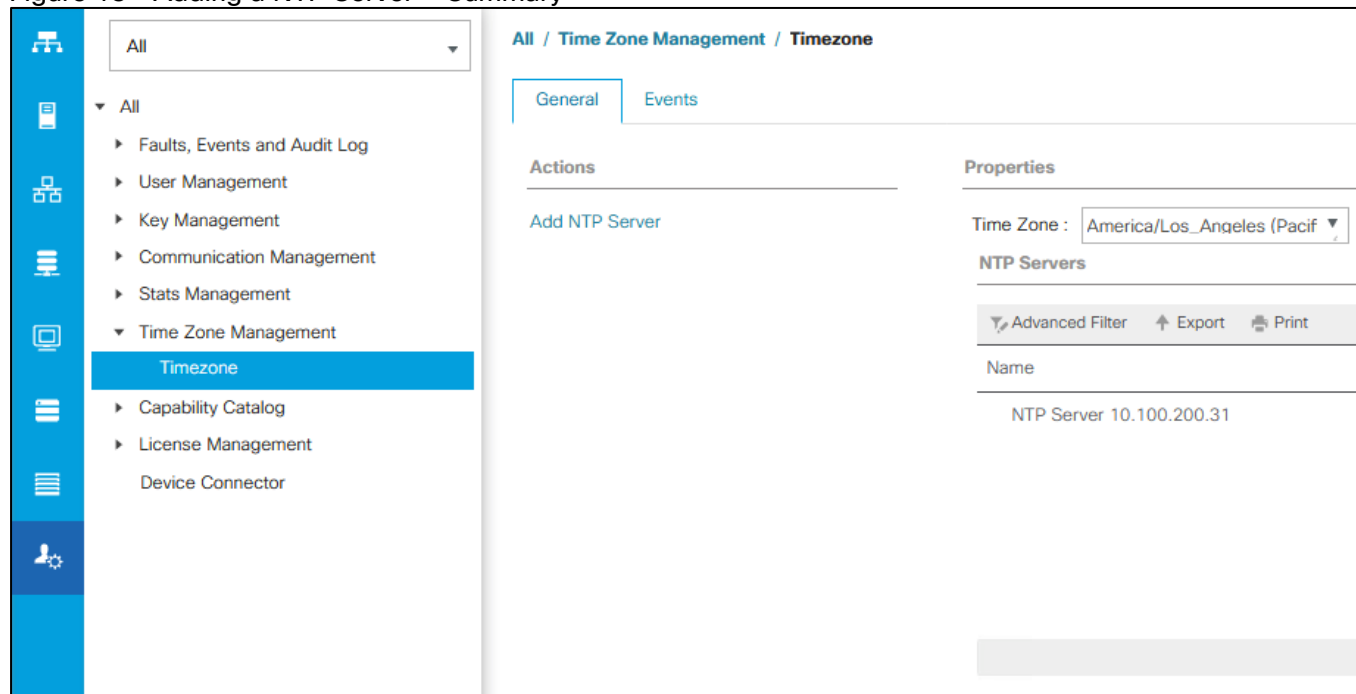
1. Open a Web browser and navigate to the Cisco UCS 6332 Fabric Interconnect cluster address.
2. Click the Launch link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. Click Launch UCS Manager HTML.
5. When prompted, enter `admin` for the username and enter the administrative password.
6. Click Login to log into the Cisco UCS Manager.

Configure NTP Server

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

1. Select Admin tab on the left side.
2. Select Time Zone Management.
3. Select Time Zone.
4. Under Properties select your time zone.
5. Select Add NTP Server.
6. Enter the IP address of the NTP server.
7. Select OK.

Figure 13 Adding a NTP Server – Summary



Initial Base Setup of the Environment

Configure Global Policies

To configure the Global Policies, follow these steps:

1. Select the Equipment tab on the left site of the window.
2. Select Policies on the right site.
3. Select Global Policies.
4. Under Chassis/FEX Discovery Policy select Platform Max under Action.
5. Select 40G under Backplane Speed Preference.
6. Under Rack Server Discovery Policy select Immediate under Action.
7. Under Rack Management Connection Policy select ~~Auto~~ Acknowledged under Action.
8. Under Power Policy select Redundancy N+1.
9. Under Global Power Allocation Policy select Policy Driven.
10. Select Save Changes.

Figure 14 Configuration of Global Policies

Equipment

Main Topology View Fabric Interconnects Servers Thermal Decommissioned Firmware Management **Policies** Faults

Global Policies Autoconfig Policies Server Inheritance Policies Server Discovery Policies SEL Policy Power Groups Port A

Chassis/FEX Discovery Policy

Action : Platform Max

Link Grouping Preference : ☒ None ☐ Port Channel

Backplane Speed Preference : ☒ 40G ☐ 4x10G

Rack Server Discovery Policy

Action : ☒ Immediate ☐ User Acknowledged

Scrub Policy : <not set>

Rack Management Connection Policy

Action : ☒ Auto Acknowledged ☐ User Acknowledged

Power Policy

Redundancy : ☐ Non Redundant ☒ N+1 ☐ Grid

MAC Address Table Aging

Aging Time : ☐ Never ☒ Mode Default ☐ other

Global Power Allocation Policy

Allocation Method : ☐ Manual Blade Level Cap ☒ Policy Driven Chassis Group Cap

Firmware Auto Sync Server Policy

Sync State : ☒ No Actions ☐ User Acknowledge

Info Policy

Action : ☒ Disabled ☐ Enabled

Global Power Profiling Policy

Profile Power : ☐

Hardware Change Discovery Policy

Action : ☒ User Acknowledged ☐ Auto Acknowledged

Enable Fabric Interconnect A Ports for Server

To enable server ports, follow these steps:

1. Select the Equipment tab on the left site.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (subordinate) > Fixed Module.
3. Click Ethernet Ports section.
4. Select Ports 1-3, right-click and then select Configure as Server Port.
5. Click Yes and then OK.
6. Repeat the same steps for Fabric Interconnect B.

Enable Fabric Interconnect A Ports for Uplinks

To enable uplink ports, follow these steps:

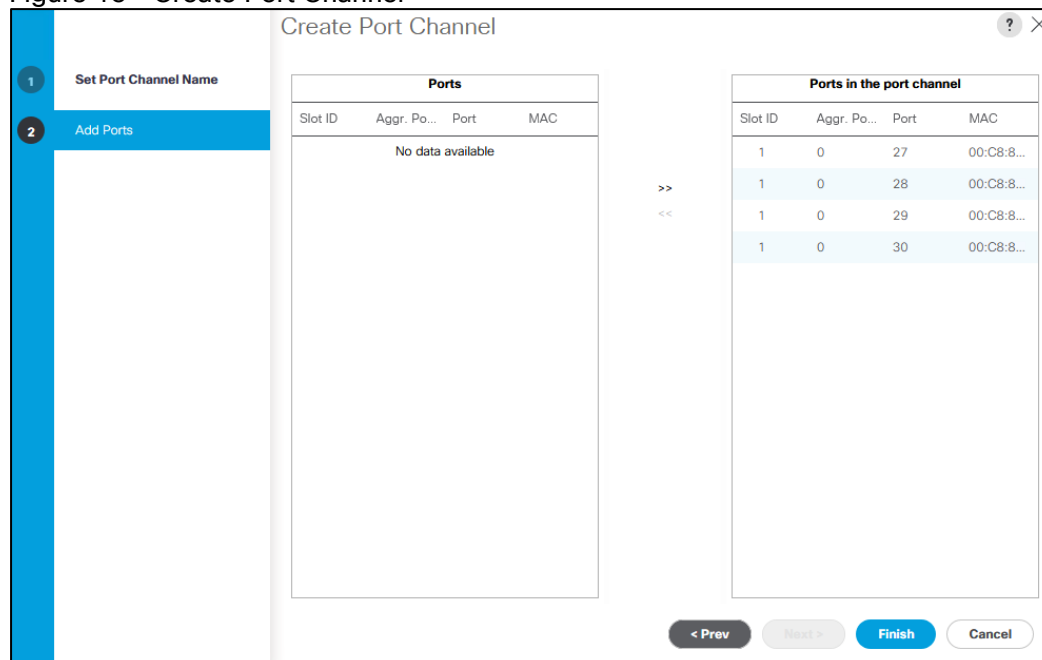
1. Select the Equipment tab on the left site.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (subordinate) > Fixed Module.
3. Click Ethernet Ports section.
4. Select Ports 27-30, right-click and then select Configure as Uplink Port.
5. Click Yes and then OK.
6. Repeat the same steps for Fabric Interconnect B.

Create Port Channel for Fabric Interconnect A/B

To create Port Channels to the connected Nexus 9332PQ switches, follow these steps:

1. Select the LAN tab in the left pane of the Cisco UCS Manager GUI.
2. Go to LAN > LAN Cloud > Fabric A > Port Channels and right-click Create Port Channel.
3. Type in ID 10.
4. Type in vPC10 in the Name field.
5. Click Next.
6. Select the available ports on the left 27-30 and assign them with >> to Ports in the Port Channel.

Figure 15 Create Port Channel



7. Click Finish and then OK.
8. Repeat the same steps for Fabric B under LAN > LAN Cloud > Fabric B > Port Channels and right-click Create Port Channel.
9. Type in ID 11.
10. Type in vPC11 in the Name field.
11. Click Next.
12. Select the available ports on the left 27–30 and assign them with >> to Ports in the Port Channel.
13. Click Finish and then OK.

Label Each Server for Identification

To label each chassis for better identification, follow these steps:

1. Select the Equipment tab on the left site.
2. Select Rack-Mounts > Servers > Server 1.
3. In the Properties section on the right go to User Label and add slicestor 1 to the field.
4. Repeat the previous steps for Server 2 – 3 by using the following labels (Table 6):

Table 6 Server Label

Server	Name
Server 1	slicestor 1
Server 2	slicestor 2
Server 3	slicestor 3

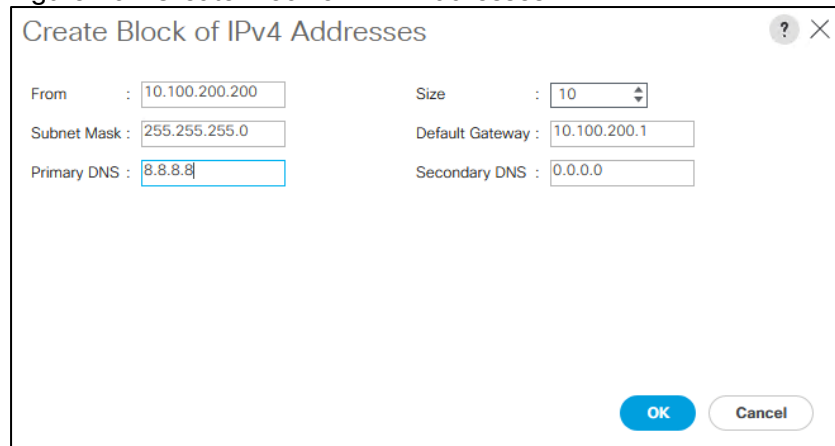
Create KVM IP Pool

To create a KVM IP Pool, follow these steps:

1. Select the LAN tab on the left site.
2. Go to LAN > Pools > root > IP Pools and right-click Create Block of IPv4 Addresses.
3. Type in IBM-IP as Name.
4. (Optional) Enter a Description of the MAC Pool.
5. Set Assignment Order as Sequential.
6. Click Next and then Add.
7. Enter an IP Address in the From field.

8. Enter **Size** 10.
9. Enter your Subnet Mask.
10. Fill in your Default Gateway.
11. Enter your Primary DNS and Secondary DNS if needed.
12. Click **OK**.

Figure 16 Create Block of IPv4 Addresses



The screenshot shows a dialog box titled "Create Block of IPv4 Addresses". It has a standard Windows-style title bar with a question mark icon and a close button (X). The dialog contains the following fields and values:

Field	Value
From	10.100.200.200
Size	10
Subnet Mask	255.255.255.0
Default Gateway	10.100.200.1
Primary DNS	8.8.8.8
Secondary DNS	0.0.0.0

At the bottom right of the dialog are two buttons: "OK" (highlighted in blue) and "Cancel".

Create MAC Pool

To create a MAC Pool, follow these steps:

1. Select the LAN tab on the left site.
2. Go to LAN > Pools > root > Mac Pools and right-click Create MAC Pool.
3. Type in **IBM-MAC** as Name.
4. (Optional) Enter a Description of the MAC Pool.
5. Set Assignment Order as Sequential.
6. Click Next.
7. Click Add.
8. Specify a starting MAC address.
9. Specify a size of the MAC address pool, which is sufficient to support the available server resources, for example, 100.

Figure 17 Create a Block of MAC Addresses

10. Click OK.

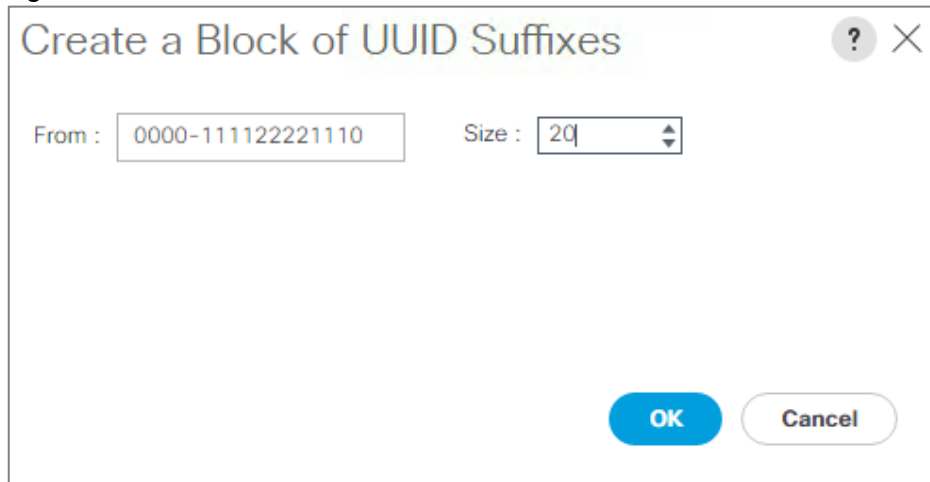
11. Click Finish.

Create UUID Pool

To create a UUID Pool, follow these steps:

1. Select the Servers tab on the left site.
2. Go to Servers > Pools > root > UUID Suffix Pools and right-click Create UUID Suffix Pool.
3. Type in **IBM-UUID** as Name.
4. (Optional) Enter a Description of the UUID Pool.
5. Set Assignment Order to Sequential and click Next.
6. Click Add.
7. Specify a starting UUID Suffix.
8. Specify a size of the UUID suffix pool, which is sufficient to support the available server resources, for example, 20.

Figure 18 Create a Block of UUID Suffixes



The screenshot shows a dialog box titled "Create a Block of UUID Suffixes". It has a standard Windows-style title bar with a question mark icon and a close button (X). The main area contains two input fields: "From :" with the value "0000-111122221110" and "Size :" with the value "20". At the bottom right, there are two buttons: "OK" (blue) and "Cancel" (white with a grey border).

9. Click OK.
10. Click Finish and then OK.

Enable CDP

To enable Network Control Policies, follow these steps:

1. Select the LAN tab in the left pane of the Cisco UCS Manager GUI.
2. Go to LAN > Policies > root > Network Control Policies and right-click Create Network-Control Policy.
3. Type in **IBM-CDP** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Click Enabled under CDP.
6. Click Only Native VLa under MAC Register Mode.
7. Leave everything else untouched and click OK.
8. Click OK.

Figure 19 Create Network Control Policy

Create Network Control Policy

Name : IBM-CDP

Description :

CDP : ☐ Disabled ☒ Enabled

MAC Register Mode : ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail : ☒ Link Down ☐ Warning

MAC Security

Forge : ☒ Allow ☐ Deny

LLDP

OK Cancel

QoS System Class

To create a Quality of Service System Class, follow these steps:

1. Select the LAN tab in the left pane of the Cisco UCS Manager GUI.
2. Go to LAN > LAN Cloud > QoS System Class.
3. Set Weight to none for Platinum, Gold, Silver, Bronze, and Fibre Channel.
4. Set Best Effort Weight to Best Effort and MTU to 9216.
5. Set Fibre Channel Weight to None.
6. Click Save Changes and then OK.

Figure 20 QoS System Class

LAN / LAN Cloud / QoS System Class

General Events FSM

Actions Properties

Use Global Owner : Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	none	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	none	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	none	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	none	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	best-effort	100	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	none	N/A	fc	N/A

vNIC Template Setup

The next step is to create the appropriate vNIC template. For IBM COS we need to create one vNIC. This vNIC will handle Management, Data, and Client traffic.

To create the appropriate vNIC, follow these steps:

1. Select the LAN tab in the left pane of the Cisco UCS Manager GUI.
2. Go to LAN > Policies > root > vNIC Templates and right-click Create vNIC Template.
3. Type in **IBM-COS** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Click Fabric A as Fabric ID and enable failover.
6. Click Updating Template as Template Type.
7. Select **Native-VLAN** as VLANs and click **Native VLAN**.
8. Type in 9000 for MTU Size.
9. Select **IBM-MAC** as MAC Pool.
10. Select **IBM-CDP** as Network Control Policy.
11. Click OK and then OK.

Figure 21 Setup vNIC Template for vNIC IBM-COS

Create vNIC Template

Name : IBM-COS

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"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input type="checkbox"/>	Public-VLAN	<input type="radio"/>

OK Cancel

Adapter Policy Setup

To create a specific adapter policy for ClevOS, follow these steps:

1. Select the Server tab in the left pane of the Cisco UCS Manager GUI.
2. Go to Servers > Policies > root > Adapter Policies and right-click Create Ethernet Adapter Policy.
3. Type in **IBM** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Under Resources type in the following values:
 - a. Transmit Queues: 8
 - b. Ring Size: 4096
 - c. Receive Queues: 8

- d. Ring Size: 4096
 - e. Completion Queues: 16
 - f. Interrupts: 32
6. Under Options enable Receive Side Scaling (RSS).
 7. Click OK and then OK.

Figure 22 Adapter Policy for IBM

Create Ethernet Adapter Policy

Name : IBM

Description :

Resources

Pooled : ☒ Disabled ☐ Enabled

Transmit Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

Receive Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

Completion Queues : 16 [1-2000]

Interrupts : 32 [1-1024]

Options

Transmit Checksum Offload : ☐ Disabled ☒ Enabled

Receive Checksum Offload : ☐ Disabled ☒ Enabled

TCP Segmentation Offload : ☐ Disabled ☒ Enabled

TCP Large Receive Offload : ☐ Disabled ☒ Enabled

Receive Side Scaling (RSS) : ☐ Disabled ☒ Enabled

Accelerated Receive Flow Steering : ☒ Disabled ☐ Enabled

Network Virtualization using Generic Routing Encapsulation : ☒ Disabled ☐ Enabled

OK Cancel

Boot Policy Setup

To create a Boot Policy, follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > Boot Policies and right-click Create Boot Policy.
3. Type in **IBM-Boot** in the Name field.
4. (Optional) Enter a description in the Description field.

5. Click Local Devices > Add Local LUN
6. Click OK.
7. Click CIMC Mounted vMedia > Add CIMC Mounted CD/DVD
8. Click OK.
9. Click OK.

Figure 23 Create Boot Policy

?

×

Create Boot Policy

Name

:

IBM-Boot

Description

:

Reboot on Boot Order Change

:

☐

Enforce vNIC/vHBA/iSCSI Name

:

☒

Boot Mode

:

☒ Legacy
 ☐ Uefi

WARNINGS:

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

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

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

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

⊕ Local Devices

⊕ CIMC Mounted vMedia

⊕ vNICs

⊕ vHBAs

⊕ iSCSI vNICs

⊕ EFI Shell

Boot Order

+

-

⚙ Advanced Filter

↑ Export

🖨 Print

⚙

Name	Or... ▲	vNIC/...	Type	LUN ...	WWN	Slot N...	Boot ...	Boot ...	Descri...
Local LUN	1								
CIMC Mounted CD/DVD	2								

⬆ Move Up

⬇ Move Down

🗑 Delete

Set All Boot Parameters

OK

Cancel

Create Maintenance Policy Setup

To setup a Maintenance Policy, follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > Maintenance Policies and right-click Create Maintenance Policy.
3. Type in `IBM-Maint` in the Name field.

4. (Optional) Enter a description in the Description field.
5. Click User Ack under Storage Config. Deployment Policy.
6. Click User Ack under Reboot Policy.
7. Click OK and then OK.

Figure 24 Create Maintenance Policy

Create Maintenance Policy

Name : IBM-Maint

Description :

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : ☐ Immediate ☒ User Ack

Reboot Policy : ☐ Immediate ☒ User Ack ☐ Timer Automatic

☐ On Next Boot (Apply pending changes at next reboot.)

OK Cancel

Create Power Control Policy Setup

To create a Power Control Policy, follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > Power Control Policies and right-click Create Power Control Policy.
3. Type in **IBM-Power** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Click No Cap.
6. Click OK and then OK.

Figure 25 Create Power Control Policy

Create Power Control Policy

Name : IBM-Power

Description :

Fan Speed Policy : Any

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

☒ No Cap ☐ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

Create Disk Scrub Policy

To prevent failures during re-deployment of a IBM COS environment, implement a Disk Scrub Policy that is enabled when removing a profile from a server.

To create a Disk Scrub Policy, follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > Scrub Policies and right-click Create Scrub Policy.
3. Type in **IBM-Scrub** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Select Disk Scrub radio button to Yes.
6. Click OK and then click OK again.

Figure 26 Create a Disk Scrub Policy

Create Scrub Policy

Name : IBM-Scrub

Description :

Disk Scrub : ☐ No ☒ Yes

BIOS Settings Scrub : ☒ No ☐ Yes

FlexFlash Scrub : ☒ No ☐ Yes

OK Cancel

Create Host Firmware Package

To create a Host Firmware Policy, follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > Host Firmware Packages and right-click Create Host Firmware Package.
3. Type in **IBM-FW** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Under Rack Package select **4.0 (1a) C**.
6. Click OK and then click OK again.

Figure 27 Create Host Firmware Policy

Create Host Firmware Package

Name : IBM-FW

Description :

How would you like to configure the Host Firmware Package?

☒ Simple ☐ Advanced

Blade Package : <not set>

Rack Package : 4.0(1a)C

Service Pack : <not set>

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

Excluded Components:

- ☐ Adapter
- ☐ BIOS
- ☐ Board Controller
- ☐ CIMC
- ☐ FC Adapters
- ☐ Flex Flash Controller
- ☐ GPUs
- ☐ HBA Option ROM
- ☐ Host NIC
- ☐ Host NIC Option ROM
- ☒ Local Disk
- ☐ NVME Mswitch Firmware
- ☐ PSU
- ☐ SAS Expander

OK Cancel

Create vMedia Policy in Cisco UCS Manager

To simplify the installation of the hardware agnostic IBM image, create the vMedia policy for the IBM Service Profile and follow these steps:

1. Select the Servers tab in the left pane.
2. Go to Servers > Policies > root > vMedia Policies and right-click Create vMedia Policy.
3. Type in **IBM-vMedia** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Click Add.
6. Type in **IBM-ISO** in the Name field.
7. (Optional) Enter a description in the Description field.
8. Click **CDD** for Device Type.
9. Click **NFS** for Protocol.
10. Type in the Hostname/IP Address **10.100.201.39**.
11. Type in **clevos-3.13.6.33-allinone-usbiso.iso** for Remote File.

12. Type in `/mnt/ibm` for Remote Path.

13. Click OK.

Figure 28 Create vMedia Mount for IBM COS Boot Image

Properties for: IBM-ISO

General Events

Actions

Delete

Properties

Name : **IBM-ISO**

Description :

Device Type : ☒ CDD ☐ HDD

Protocol : ☒ NFS ☐ CIFS ☐ HTTP ☐ HTTPS

Hostname/IP Address : 10.100.201.39

Image Name Variable : ☒ None ☐ Service Profile Name

Remote File : clevos-3.13.6.33-allinone-usbiso.iso

Remote Path : /mnt/ibm

Remap on Eject : ☐

OK Apply Cancel Help

Creating Storage Profiles

In the next part we're going to create the Disk Group Policy and Storage Profile for the boot devices for the rear end SATA SSDs.

Creating Disk Group Policy for Boot Devices

To create the Disk Group Policy from the rear end SATA SSDs, follow these steps:

1. Select Storage in the left pane of the Cisco UCS Manager GUI.
2. Go to Storage > Storage Policies > root > Disk Group Policies and right-click Create Disk Group Policy
3. Type in **IBM-Boot** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Select **RAID 1 Mirrored** for RAID Level.
6. Click Disk Group Configuration (Manual).
7. Click Add.

8. Type in 13 as slot number.
9. Repeat the step for slot number 14.
10. Leave everything as default.
11. Click OK and then click OK again.

Figure 29 Create Disk Group Policy for Boot Device

Create Disk Group Policy

Name :

Description :

RAID Level :

☐ Disk Group Configuration (Automatic) ☒ Disk Group Configuration (Manual)

Disk Group Configuration (Manual)

Advanced Filter Export Print

Slot Number	Role	Span ID
13	Normal	Unspecified
14	Normal	Unspecified

+ Add - Delete i Info

Virtual Drive Configuration

Strip Size (KB) :

Access Policy : ☒ Platform Default ☐ Read Write ☐ Read Only ☐ Blocked

OK Cancel

Create Storage Profile

To create the Storage Profile, follow these steps:

1. Select Storage in the left pane of the Cisco UCS Manager GUI.
2. Go to Storage > Storage Profiles and right-click Create Storage Profile.
3. Type in **IBM-COS** in the Name field.
4. (Optional) Enter a description in the Description field.
5. Click Add under Local LUN.
6. Type in **Boot** in the Name field.
7. Click Expand To Available.
8. Select **IBM-Boot** under Select Disk Group Configuration.
9. Click OK, then click OK, and click OK again.

Create Service Profile Template

Create Service Profile Template

To create the Service Profile Template, follow these steps:

1. Select Servers in the left pane of the Cisco UCS Manager GUI.
2. Go to Servers > Service Profile Templates and right-click Create Service Profile Template.

Identify Service Profile Template

1. Type in **IBM** in the Name field.
2. Click Updating Template in Type.
3. In the UUID Assignment section, select the **IBM-UUID** Pool.
4. (Optional) Enter a description in the Description field.
5. Click Next.

Storage Provisioning

1. Go to the Storage Profile Policy tab and select the Storage Profile **IBM-COS**.
2. Click Next.

Networking

1. Select the Expert radio button for the option How would you like to configure LAN connectivity?
2. Click Add to add a vNIC to the template.
3. Insert **IBM** as Name.
4. Select Use vNIC Template.
5. Select **IBM-COS** under vNIC Template.
6. Select **IBM** as Adapter Policy.
7. Click OK.
8. Click Next to continue with SAN Connectivity.
9. Select No vHBA for How would you like to configure SAN Connectivity?
10. Click Next to continue with Zoning.
11. Click Next to continue with vNIC/vHBA Placement.

12. Click Next to continue with vMedia Policy.

vMedia Policy

1. Select **IBM-vMedia** from the vMedia Policy Menu.
2. Click Next.

Server Boot Order

1. Select **IBM-Boot** from the Boot Policy Menu.
2. Click Next.

Server Maintenance

1. Select the Maintenance Policy **IBM-Maint** under Maintenance Policy.
2. Click Next.
3. Click Next.

Server Assignment

1. Under Firmware Management select **IBM-FW**.
2. Click Next.

Operational Policies

1. Under Power Control Policy Configuration select **IBM-Power**, under Scrub Policy select **IBM-Scrub**, and under Management IP Address select **IBM-IP** for Outband IPv4.
2. Click Finish.

Create Service Profiles from Template

Create the appropriate Service Profiles from the previous Service Profile Template. To create all three profiles for the IBM Server, follow these steps:

1. Select Servers from the left pane of the Cisco UCS Manager GUI.
2. Go to Servers > Service Profiles and right-click Create Service Profiles from Template.
3. Type in **slicestor** in the Name Prefix field.
4. Type 1 for Name Suffix Starting Number.
5. Type 3 for Number of Instances.
6. Choose **IBM** under Service Profile Template.

7. Click OK.

Associate Service Profiles

To associate the service profiles, follow these steps:

1. Right-click the service profile `s1icestor1` and choose Change Service Profile Association.
2. Server Assignment should be Select Existing Server.
3. Select Rack ID 1.
4. Click OK and Yes and OK.
5. Repeat the steps for `s1icestor2` and Rack ID 2. Repeat the steps for the third Service Profiles counting up the Rack ID number corresponding with the service profile.

The formal setup of the Cisco UCS Manager environment and both Cisco Nexus 9332PQ switches is now finished and the installation of the IBM Cloud Object Storage ClevOS software will continue.

Installation of IBM Cloud Object Storage

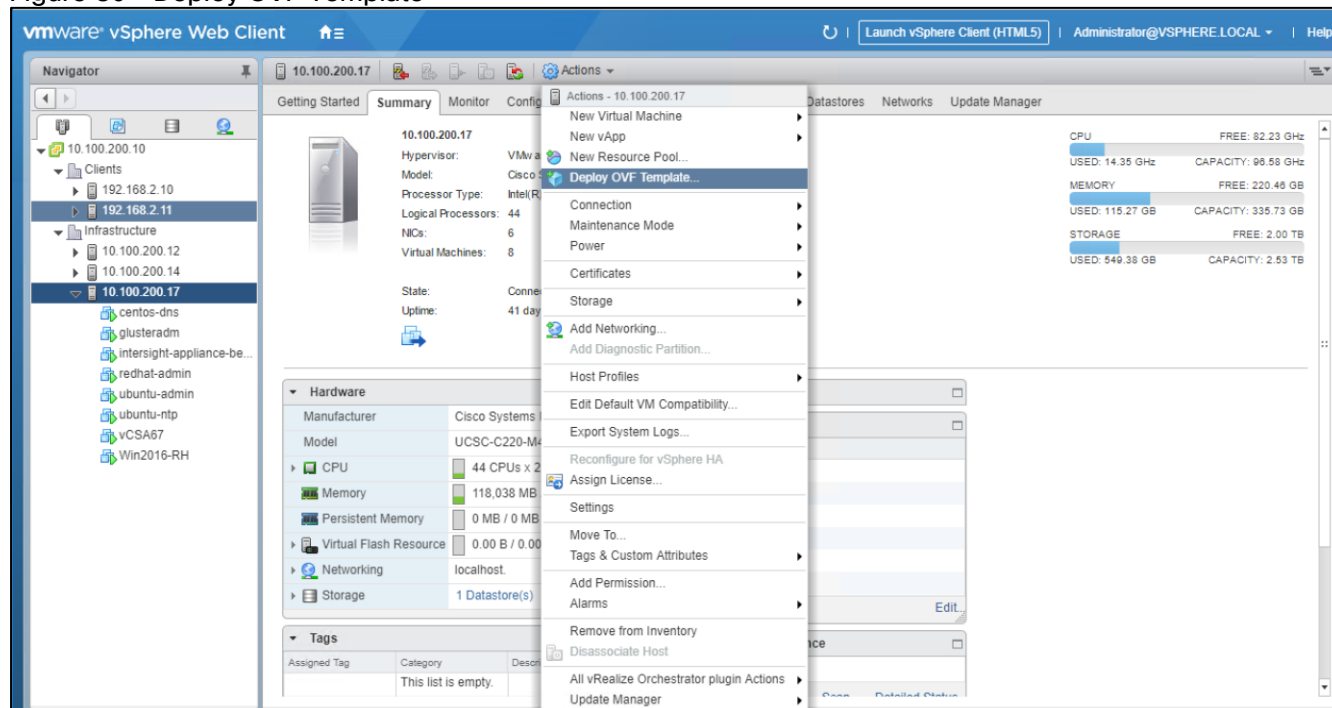
The following section provides detailed information about the installation of vManager for IBM COS on VMware vCenter and Slicestors with embedded Accessers on Cisco UCS C240 M5L.

Deployment of Virtual IBM COS Manager on VMware vCenter

To deploy the virtual IBM COS Manager on VMware vCenter, follow these steps:

1. Log into your local vCenter and select the ESXi host you want to use for deploying the virtual appliance.
2. Select under Actions – Deploy OVF Template

Figure 30 Deploy OVF Template



3. Select the template clevos-3.13.6.33-manager.ova on your local filesystem and then click Next to select your datacenter.

Figure 31 Select the Datacenter for the Location of the Appliance

The screenshot shows the 'Deploy OVF Template' wizard at step 2, 'Select name and location'. The left sidebar shows the progress: 1. Select template (checked), 2. Select name and location (active), 3. Select a resource, 4. Review details, 5. Select storage, and 6. Ready to complete. The main area is titled 'Select name and location' with the instruction 'Enter a name for the OVF and select a deployment location.' There is a text field for 'Name' containing 'clevos-3.13.6.33-manager'. Below it are 'Filter' and 'Browse' buttons. The instruction 'Select a datacenter or folder.' is followed by a tree view showing a folder '10.100.200.10' expanded, with sub-items 'Clients' and 'Infrastructure'. 'Infrastructure' is selected and highlighted in blue. At the bottom right are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

4. Click Next to choose your resource and then click Next

Figure 32 Select the Resource for the Appliance

The screenshot shows the 'Deploy OVF Template' wizard at step 3, 'Select a resource'. The left sidebar shows the progress: 1. Select template (checked), 2. Select name and location (checked), 3. Select a resource (active), 4. Review details, 5. Select storage, and 6. Ready to complete. The main area is titled 'Select a resource' with the instruction 'Select where to run the deployed template.' There are 'Filter' and 'Browse' buttons. The instruction 'Select a host, cluster, resource pool or vapp.' is followed by a tree view showing a folder 'Infrastructure' expanded, with sub-items '10.100.200.12', '10.100.200.14', and '10.100.200.17'. '10.100.200.17' is selected and highlighted in blue. At the bottom right are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

- Review the details and click Next to select the storage.

Figure 33 Select the Storage

Deploy OVF Template

1 Select template
2 Select name and location
3 Select a resource
4 Review details
5 Select storage
6 Select networks
7 Ready to complete

Select storage
Select location to store the files for the deployed template.

Select virtual disk format: **Thick provision lazy zeroed**

VM storage policy: **None**

☐ Show datastores from Storage DRS clusters ⓘ

Filter

Datastores **Datastore Clusters**

Name	Status	VM storage policy	Capacity	Free
C220_Flash_ESXi	✓ Normal	VM Encryption P...	2.53 TB	2 TB

1 Objects Copy

Back **Next** **Finish** **Cancel**

- Click Next and select the network.

Figure 34 Select the Network to use for the Virtual Appliance

Deploy OVF Template

- 1 Select template
- 2 Select name and location
- 3 Select a resource
- 4 Review details
- 5 Select storage
- 6 Select networks**
- 7 Ready to complete

Select networks
Select a destination network for each source network.

Source Network	Destination Network
Network 1	VM Network 192.168.2.X


Description - Network 1
Network 1

IP Allocation Settings
IP protocol: IPv4
IP allocation: Static - Manual ⓘ

Back Next Finish Cancel

7. Click Next, review the summary and then click Finish to deploy the virtual appliance.
8. When the deployment finished, please start the virtual machine and launch the remote console.

Figure 35 Launching the Remote Console



9. At the ClevOS Manager login prompt, provide the following credentials:
 - a. Username – localadmin
 - b. Password – password
10. When logged in at the local console, change the password by following the prompts:
 - a. manager# password
 - b. Current password:
 - i. New password: <type in the new, secure password here>
 - ii. Retype new password: <re-enter the new password from above here>
 - iii. Password change successful
 - iv. manager#



It is highly recommended to change the password at first login. In addition to following good security pro-tocol, ClevOS will not enable Secure Shell (SSH) remote access until the default password has been changed.

11. ClevOS uses a configuration shell that can be entered by entering the command edit. Enter the configuration shell and input the following commands to perform initial configuration steps, making changes in your environment as necessary:

```
manager# edit
```



```
manager (working)#
```

Check the network interface

```
manager (working)# port
```

PORT	ADDRESS	MAX SPEED	STATUS
eth0	00:50:56:ba:3a:62	10000 Mbps	disconnected

Configure the interface that will be part of the channel data

```
manager (working)# channel data port eth0
```

Establish an IP address for the data channel

```
manager (working)# channel data ip 192.168.2.100 netmask 255.255.255.0 gateway 192.168.2.1
```

Configure the hostname for the appliance

```
manager (working)# system hostname manager
```

Provide some basic location details. This increases the randomness during the private key generation process.

```
manager (working)# system organization cisco city sjc state ca country us
```

Once all information has been entered suitable to the deployment, activate the configuration.

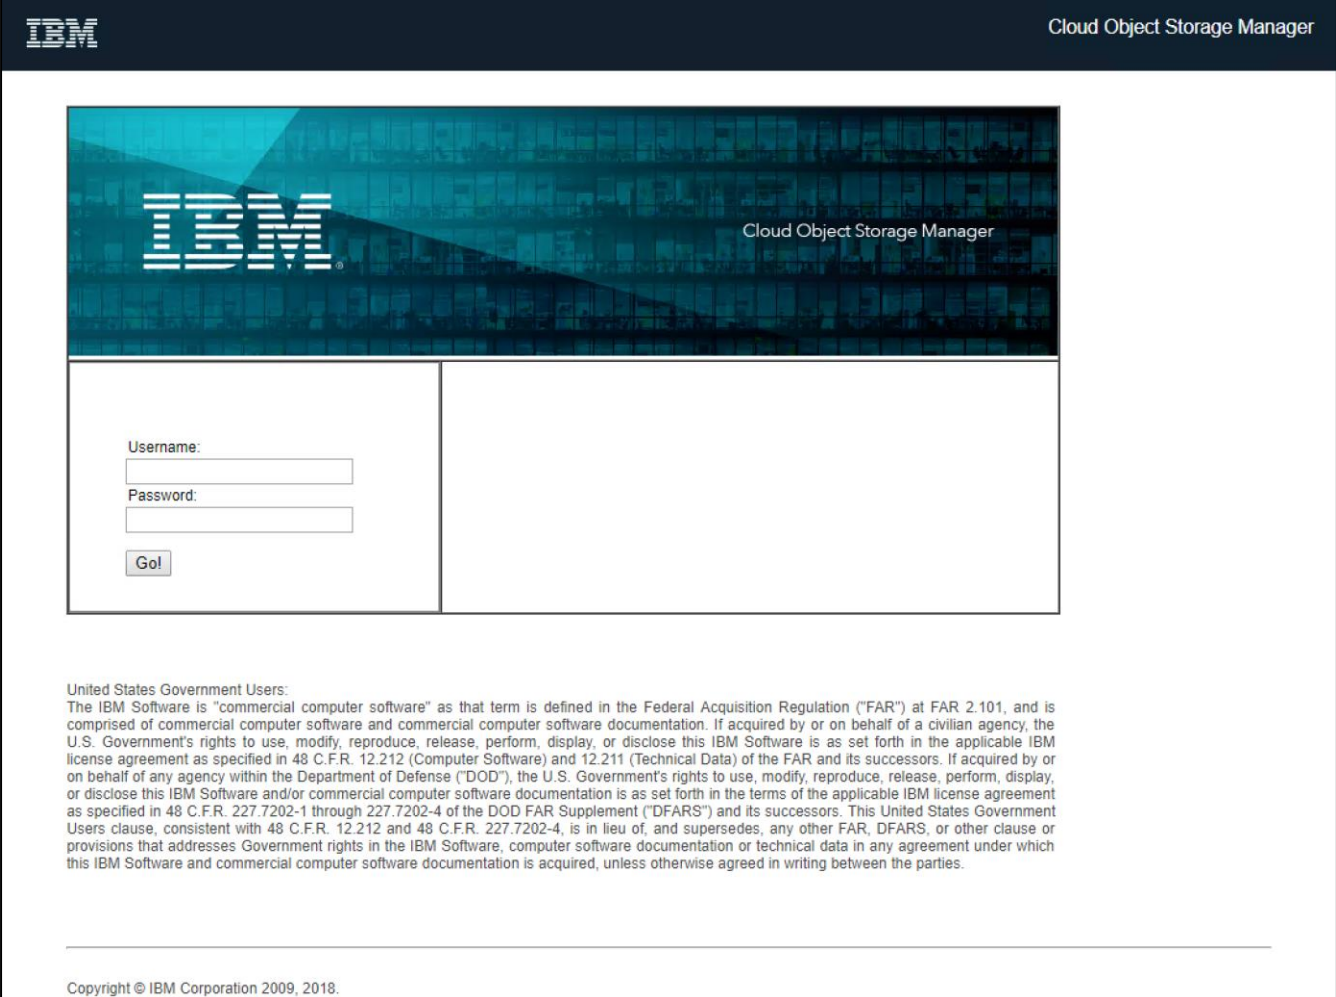
```
manager (working)# activate
```

Wait for activation to complete. The servers in this design guide exist on a private network; therefore, any DNS or gateway errors are safe to ignore.

Once activation has completed, navigate in a browser to the ClevOS Manager ip address configured up above 192.168.2.100.

If the webserver responds and the following screen appears, initial first time console setup has completed.

Figure 36 Log in screen for Cloud Object Storage Manager



IBM Cloud Object Storage Manager

IBM Cloud Object Storage Manager

Username:

Password:

Go!

United States Government Users:
 The IBM Software is "commercial computer software" as that term is defined in the Federal Acquisition Regulation ("FAR") at FAR 2.101, and is comprised of commercial computer software and commercial computer software documentation. If acquired by or on behalf of a civilian agency, the U.S. Government's rights to use, modify, reproduce, release, perform, display, or disclose this IBM Software is as set forth in the applicable IBM license agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the FAR and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government's rights to use, modify, reproduce, release, perform, display, or disclose this IBM Software and/or commercial computer software documentation is as set forth in the terms of the applicable IBM license agreement as specified in 48 C.F.R. 227.7202-1 through 227.7202-4 of the DOD FAR Supplement ("DFARS") and its successors. This United States Government Users clause, consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-4, is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provisions that addresses Government rights in the IBM Software, computer software documentation or technical data in any agreement under which this IBM Software and commercial computer software documentation is acquired, unless otherwise agreed in writing between the parties.

Copyright © IBM Corporation 2009, 2018.

12. First time console configuration should now be complete on the ClevOS Manager Node. To finish initial web UI configuration, complete the following steps:
 - a. At the IBM Cloud Object Storage Manager login prompt at the IP address configured above, provide the following credentials:
 - i. Username – admin
 - ii. Password – password
13. When logged in web interface, accept the license agreement and input the name of the license acceptor. Select the button titled **Accept IBM & non-IBM Licenses**.

Figure 37 Accept IBM & non-IBM Licenses

IBM Admin | Help | Sign Out
Cloud Object Storage Manager

End User License Agreement Print Decline Accept IBM & non-IBM Licenses

Please read and accept the following IBM and non-IBM license agreements before proceeding. Language: English

LICENSE INFORMATION

The Programs listed below are licensed under the following License Information terms and conditions in addition to the Program license terms previously agreed to by Client and IBM. If Client does not have previously agreed to license terms in effect for the Program, the International Program License Agreement (Z125-3381-14) applies.

Program Name (Program Number):

- IBM Cloud Object Storage System 3.13.6 (S725-Z81)
- IBM Cloud Object Storage 1YR 3.13.6 (S641-C01)
- IBM Cloud Object Storage 2YR 3.13.6 (S641-C02)
- IBM Cloud Object Storage 3YR 3.13.6 (S641-C03)
- IBM Cloud Object Storage 4YR 3.13.6 (S641-C04)
- IBM Cloud Object Storage 5YR 3.13.6 (S641-C05)
- IBM Cloud Object Storage System FIPS 3.13.6 (S725-Z81)
- IBM Cloud Object Storage FIPS 1YR 3.13.6 (S641-C01)
- IBM Cloud Object Storage FIPS 2YR 3.13.6 (S641-C02)
- IBM Cloud Object Storage FIPS 3YR 3.13.6 (S641-C03)
- IBM Cloud Object Storage FIPS 4YR 3.13.6 (S641-C04)
- IBM Cloud Object Storage FIPS 5YR 3.13.6 (S641-C05)

The following standard terms apply to Licensee's use of the Program.

Limited use right

As described in the International Program License Agreement ("IPLA") and this License Information, IBM grants Licensee a limited right to use the Program. This right is limited to the level of Authorized Use, such as a Processor Value Unit ("PVU"), a Resource Value Unit ("RVU"), a Value Unit ("VU"), or other specified level of use, paid for by Licensee as evidenced in the Proof of Entitlement. Licensee's use may also be limited to a specified machine, or only as a Supporting Program, or subject to other restrictions. As Licensee has not paid for all of the economic value of the Program, no other use is permitted without the payment of additional fees. In addition, Licensee is not authorized to use the Program to provide commercial IT services to any third party, to provide commercial hosting or timesharing, or to sublicense, rent, or lease the Program unless expressly provided for in the applicable agreements under which Licensee obtains authorizations to use the Program. Additional rights may be available to Licensee subject to the payment of additional fees or under different or supplementary terms. IBM reserves the right to determine whether to make such additional rights available to

☒ IBM License Agreement
☐ Non-IBM License Agreement

☒ I have read and agreed to the terms provided in the IBM and non-IBM license agreements (required for acceptance).

Print Name (License Acceptor):

Print Decline Accept IBM & non-IBM Licenses

Copyright © 2009-2018 IBM Corporation. All rights reserved.
[IBM Trademarks and Patents](#)

page served 2018-12-17 10:35:45 GMT from 192.168.2.100
v. 3.13.6.33

14. At the next screen, select the radio button Create a new system and then select the Begin button.

Figure 38 Create new system

IBM Admin | Help | Sign Out
Cloud Object Storage Manager

Introduction Begin »

To begin, select one of the options below:

☒ Create a new system
☐ Restore this manager from a manager backup file

Begin »

Copyright © 2009-2018 IBM Corporation. All rights reserved.
[IBM Trademarks and Patents](#)

page served 2018-12-17 10:36:30 GMT from 192.168.2.100
v. 3.13.6.33

15. At the Admin Password screen, enter a new password in the both fields to change the default and then select Save and Continue.

Figure 39 Enter new password

16. At the Create Sites screen, modify as much information as desired and then click **Finish**.

Figure 40 Create sites

help.' A note indicates '* indicates required field'. There is a '+ Add Additional Site' button. The form contains several input fields: 'Site 1*:' (with 'My Site' entered), 'Abbreviation:', 'Description:', 'Company:' (with 'Cisco Systems' entered), 'Address:', 'Phone:', 'Latitude:', and 'Longitude:'. A 'Finish' button is at the bottom right of the form area. The footer contains copyright information: 'Copyright © 2009-2018 IBM Corporation. All rights reserved. IBM Trademarks and Patents' and server details: 'page served 2018-12-17 10:44:56 GMT from 192.168.2.100 v. 3.13.6.33'."/>

The initial set up of the virtual appliance Manager is now finished.

Deployment of IBM COS Slicestor on Cisco UCS C240 M5L

To install ClevOS onto the Cisco UCS C240 M5L to be used as the IBM Cloud Object Storage Slicestor, follow these steps:

1. Select the **Servers** button on the left-hand side.

2. Navigate to `Servers > Service Profiles > root > slicestor-1` from the exposed, left-hand tree.
3. Select `KVM Console` underneath the `Actions` section of the right-hand pane. Accept any prompts or follow any links until the KVM Console is present. This could require a Java software upgrade or disabling pop ups in the browser.
4. Wait until the ClevOS Installer appears and select then `#1 Perform Automatic Installation`.

Figure 41 ClevOS Installer

```
*****
IBM Cloud Object Storage System (r) Installer
*****

#1.      Perform automatic installation
#2.      Perform manual installation
#3.      Reboot
Choose action:  (1-3): _
```

5. Select option `#2 Factory Install (Erase all disks and install)` at the next ClevOS installation screen that appears.

Figure 42 Disk erase choice

```

*****
IBM Cloud Object Storage System (r) Installer
*****

#1.      OS Disk Only (Erase only OS disk and install)
#2.      Factory Install (Erase all disks and install)
Select installation type:  (1-2):

```

6. When making the last selection, a new prompt will appear warning that all disks will be erased during this process. To confirm this data destructive behavior, type in `erase` and hit enter.

Figure 43 Confirmation of disk erase

```

*****
IBM Cloud Object Storage System (r) Installer
*****

#1.      OS Disk Only (Erase only OS disk and install)
#2.      Factory Install (Erase all disks and install)
Select installation type:  (1-2): 2

WARNING:  This option will erase all disks attached to the system.
Enter 'erase' (no quotes) to confirm.  Other input will cancel:

```

7. At the next ClevOS installation screen, select the desired source image, #3 CLEVOS-3.13.6.33~SLICESTOR and hit enter.

Figure 44 Select SLICESTOR installation

```
*****
IBM Cloud Object Storage System (r) Installer
*****

#1.          CLEVOS-3.13.6.33-ACCESSER
#2.          CLEVOS-3.13.6.33-MANAGER
#3.          CLEVOS-3.13.6.33-SLICESTOR
#4.          CLEVOS-3.13.6.33-SMC
Choose source image (1-4):
```

8. When the system has been rebooted and the following screen appears, ClevOS installation has completed for node slicestor1.

Figure 45 Login screen for slicestor1

```
IBM Cloud Object Storage
Cisco UCS C240 M5 ClevOS 3.13.6.33 tty1
slicestor login:
```

9. Installation has finished on the IBM COS Slicestor node. To finish initial Slicestor node first-time configuration, complete the following steps:
 - a. At the ClevOS Slicestor login prompt, provide the following credentials:
 - i. Username - localadmin
 - ii. Password - password
10. When logged in at the local console, change the password by following the prompts:
 - a. slicestor# password
 - b. Current password:
 - i. New password: <type in the new, secure password here>
 - ii. Retype new password: <re-enter the new password from above here>
 - iii. Password change successful
 - iv. slicestor#



It is highly recommended to change the password at first login. In addition to following good security pro-tocol, ClevOS will not enable Secure Shell (SSH) remote access until the default password has been changed.

11. ClevOS uses a configuration shell that can be entered by entering the command edit. Enter the configuration shell and input the following commands to perform initial configuration steps, making changes in your environment as necessary:

```
slicestor# edit
slicestor (working)#
```


Check the network interface

```
slicestor (working)# port
```

```
PORT ADDRESS MAX SPEED STATUS
p11p1 00:25:b5:f2:00:00 10000 Mbps disconnected
p11p2 00:25:b5:f2:00:05 10000 Mbps disconnected
```

Configure the interface that will be part of the channel data

```
slicestor (working)# channel data port p11p1,p11p2
```

Establish an IP address for the data channel

```
slicestor (working)# channel data ip 192.168.2.101 netmask 255.255.255.0 gateway
192.168.2.1
```

Set the bonding type to be used by the data channel.

```
slicestor (working)# channel data bonding active-backup
```

Configure MTU 9000 for the data channel.

```
slicestor (working)# channel data bondmtu 9000
```

Configure the hostname for the appliance

```
slicestor (working)# system hostname slicestor1
```

Provide some basic location details. This increases the randomness during the private key generation process.

```
slicestor (working)# system organization cisco city sjc state ca country us
```

Provide the IP address of the previously configured ClevOS Manager. Accept any errors about manager certificates which occur as a result of the network interface not yet being up by entering y and skip entering a manager prefix by selecting the enter key at the prompt.

```
slicestor (working)# manager ip 192.168.2.100
```

Manager Certificate Data

```
=====
```

Fingerprint: 05:55:82:c2:b4:00:18:0a:b9:6e:be:28:4a:b6:8c:ce:ec:c6:97:a9

Serial No.: d5d6b25c-273e-78ca-0025-c4b7a311e942

Name: dsNet Manager CA

Organization: Cleversafe

City: Chicago

State: Illinois

Country: US

Accept this certificate? (I for more info) [y/N/i]: y

```
slicestor (working)#
```

Once all information has been entered suitable to the deployment, activate the configuration.

```
slicestor (working)# activate
```

12. Wait for activation to complete. The servers in this design guide exist on a private network; therefore, any DNS or gateway errors are safe to ignore.
13. Follow the above steps to set up the remaining Slicestors 2 and 3, taking care to modify the IP address and hostname according to the table below:

Table 7 IP Address and Hostnames

Service Profile	IP Address	Hostname
slicestor1	192.168.2.101	slicestor1
slicestor2	192.168.2.102	slicestor2
slicestor3	192.168.2.103	slicestor3

IBM COS Jumbo Frame Verification

To verify that jumbo frames are correctly implemented in the environment, ClevOS has iperf installed by default on all nodes. Before moving on to more complex activities, it can be beneficial to verify the network as operating as intended. This test will also determine if SSH was configured correctly. To test if MTU 9000 is correctly configured, complete the following steps:

1. Open a Secure Shell client of choice. PuTTY will be used in this example, but others may be used as well.
2. In the field titled **Host Name (or IP Address)**, enter `localadmin@192.168.2.101` to connect to the first Slicestor.
3. Underneath the Radio Button for **Connection type**, select **SSH**. This should automatically select the correct port 22.
4. Select the **Open** button and accept any certificates popup windows until a login prompt is present.
5. Type in the password previously configured and hit the enter key until the connected to the ClevOS shell:

```
Using username "localadmin".
```

```
localadmin@192.168.2.101's password:
```

```
IBM Cloud Object Storage
```

```
Cisco UCS C240 M5 ClevOS 3.13.6.33
```

```
IBM Cloud Object Storage Device Shell
```

```
Type '?' or 'help' to get the list of available commands.
```

```
slicestor1#
```

```
Enter the root administrator shell.
```

```

slice1stor1# su
root@slice1stor1:~#
Temporarily disable the firewall so that network throughput testing may occur.
root@slice1stor1:~# service iptables stop
[ ok ] Stopping iptables: iptables.
root@slice1stor1:~#
Start iperf in server mode.
root@slice1stor1:~# iperf -s -p 5002

```

```

-----
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
-----

```

6. Follow steps 1 through 5 above on the second Slice1stor at IP address 192.168.2.102 to arrive at a similar secure shell prompt:

```

localadmin@192.168.2.102's password:
IBM Cloud Object Storage
Cisco UCS C240 M5 ClevOS 3.13.6.33
IBM Cloud Object Storage Device Shell
Type '?' or 'help' to get the list of available commands.
slice2stor#
root@slice2stor:~# service iptables stop
[ ok ] Stopping iptables: iptables.
root@slice2stor:~#

```

7. The following command `iperf -c 192.168.2.101 -P 4 -m -p 5002` will run iperf in client mode. There are two things this will test: total throughput and MTU size. The frame size is underlined in yellow and the total throughput is underlined in red.

Figure 46 Iperf Performance Test

```

root@slice2tor2:~# iperf -c 192.168.2.101 -P 4 -m -p 5002
-----
Client connecting to 192.168.2.101, TCP port 5002
TCP window size: 325 KByte (default)
-----
[ 5] local 192.168.2.102 port 58754 connected with 192.168.2.101 port 5002
[ 4] local 192.168.2.102 port 58748 connected with 192.168.2.101 port 5002
[ 3] local 192.168.2.102 port 58750 connected with 192.168.2.101 port 5002
[ 6] local 192.168.2.102 port 58752 connected with 192.168.2.101 port 5002
[ ID] Interval      Transfer    Bandwidth
[ 5] 0.0-10.0 sec  11.0 GBytes  9.45 Gbits/sec
[ 5] MSS size 8948 bytes (MTU 8988 bytes, unknown interface)
[ 4] 0.0-10.0 sec  12.0 GBytes  10.3 Gbits/sec
[ 4] MSS size 8948 bytes (MTU 8988 bytes, unknown interface)
[ 3] 0.0-10.0 sec  11.0 GBytes  9.45 Gbits/sec
[ 3] MSS size 8948 bytes (MTU 8988 bytes, unknown interface)
[ 6] 0.0-10.0 sec  12.0 GBytes  10.3 Gbits/sec
[ 6] MSS size 8948 bytes (MTU 8988 bytes, unknown interface)
[SUM] 0.0-10.0 sec  46.1 GBytes  39.6 Gbits/sec
root@slice2tor2:~#

```

8. Perform this same test on an all ClevOS Slice2tor nodes as desired to confirm that jumbo frames is enable on all servers.
9. When testing has completed, re-enable the firewall by rebooting the host or issuing the command `service iptables start`.

```

root@slice2tor1:~# service iptables start
[ ok ] Starting iptables: iptables.

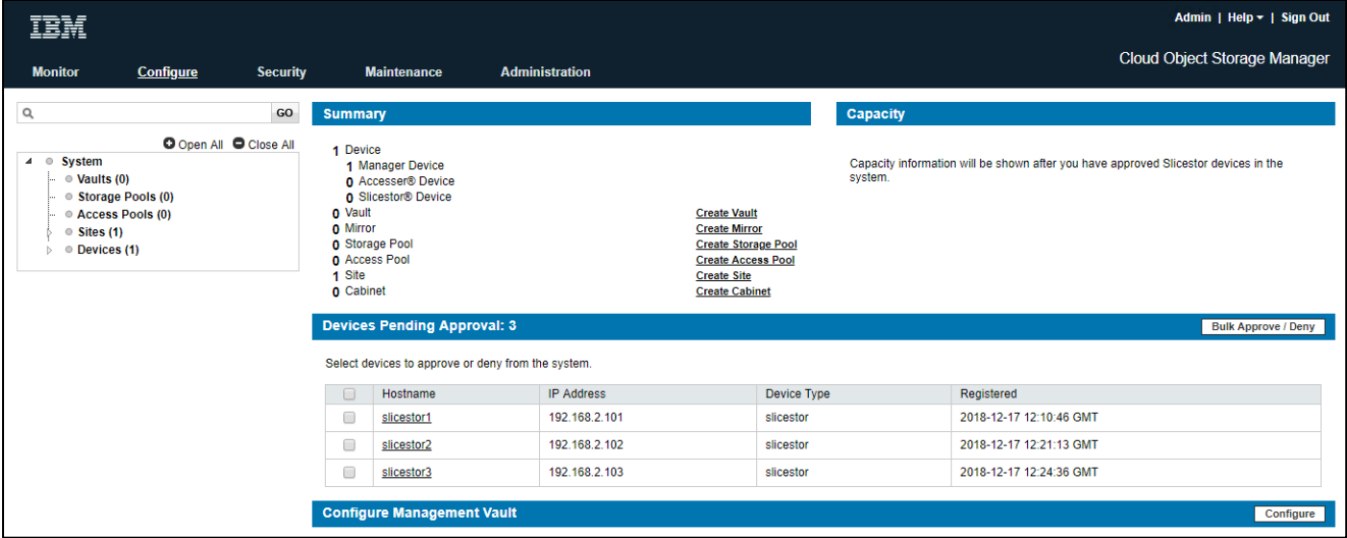
```

IBM COS dsNet Setup

All software should be installed at this point and all nodes should have contacted the Manager for entry into the IBM COS dsNet . To add all nodes to the IBM COS manager, follow these steps:

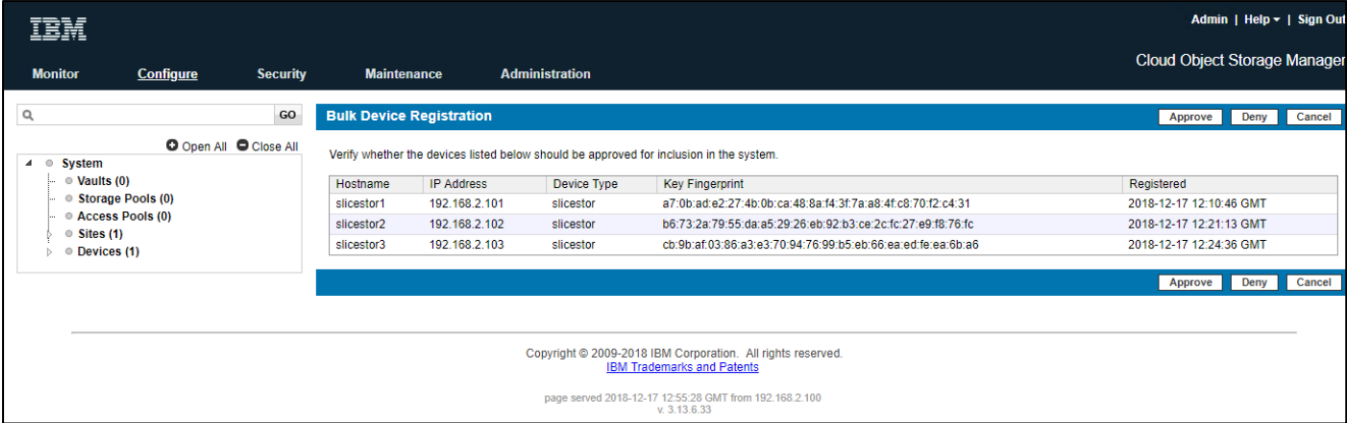
1. Open a web browser and navigate to the IP address of the IBM COS Manager 192.168.2.100 and log in with the credentials created previously.
2. Select **Configure** from the top navigation bar and observe 3 devices pending approval. This happens automatically after following the previous console command `manager ip 192.168.2.100`.

Figure 47 Bulk Approve Slicestors



3. Select the checkbox directly to the left of the column header `Hostname`. This should select all pending devices for approval. Then select the button for `Bulk Approve / Deny`.

Figure 48 Bulk Device Registration



4. At the Bulk Device Registration screen, click `Approve`.

Figure 49 Bulk Edit Device Site

IBM Cloud Object Storage Manager

Admin | Help | Sign Out

Monitor Configure Security Maintenance Administration

Bulk Edit Device Site

Select devices to assign them to a site:

<input checked="" type="checkbox"/>	Hostname
<input checked="" type="checkbox"/>	slice10r1 (192.168.2.101)
<input checked="" type="checkbox"/>	slice10r2 (192.168.2.102)
<input checked="" type="checkbox"/>	slice10r3 (192.168.2.103)

Assign the selected devices to site:

☒ Cisco Validated Design
☐ Or, create a new site for the selected devices

New Site Name:

Save

Copyright © 2009-2018 IBM Corporation. All rights reserved.
[IBM Trademarks and Patents](#)
 page served 2018-12-17 12:58:02 GMT from 192.168.2.100
 v. 3.13.6.33

- At the Bulk Edit Device Site screen, select the checkbox directly to the left of the column header Hostname. Next, select the radio button for the previously created site name, Cisco Validated Design. Finally, select the Save button at the bottom right.
- (Optional) At the Bulk Edit Device Alias screen, provide any alias beyond the hostname for each node if desired. Once complete, or if no alias required, select the Save button at the bottom right.
- Device registration should now be complete and all nodes added to the Site.

Configure IBM COS to Sync with an NTP Server

It is critical to have time in sync both across all COS nodes. In order to configure the IBM COS dsNet to sync with an NTP server, follow these steps:

- Open a web browser and navigate to the IP address of the IBM COS Manager 192.168.2.100 and log in with the credentials created previously.
- Select the Administration tab from the topmost navigation bar and scroll down to System NTP Configuration and select Configure
- Select the middle radio button next to Manager and External NTP. Next, enter the IP address of the configured NTP server in the External NTP Servers dialog box. Select the Update button at the bottom right.

Figure 50 System NTP Configuration

IBM Cloud Object Storage Manager

Admin | Help | Sign Out

Monitor Configure Security Maintenance Administration

System NTP Configuration

Configure the NTP behavior for devices in the system.

☐ Manager NTP Only
 All devices sync to the manager only. The manager syncs to the external NTP servers.

☒ Manager And External NTP
 All devices sync to both the manager and the external NTP servers.

☐ External NTP Only
 All devices sync only to the external NTP servers.

External NTP Servers:

192.168.2.31

Cancel Update

Copyright © 2009-2018 IBM Corporation. All rights reserved.
[IBM Trademarks and Patents](#)
 page served 2018-12-17 13:07:27 GMT from 192.168.2.100
 v. 3.13.6.33

Configure Access Key Authentication

Configuring IBM COS to accept an access key can provide additional security and flexibility for object storage users. To enable access key authentication, follow these steps:

1. Select the **Security** tab from the topmost navigation bar.
2. Select the **Configure** button in the section titled **Enable/Disable Authentication Mechanisms**.
3. Select the check box next to **Enable access key authentication** and then select **Update**.

Figure 51 Enable / Disable Authentication Mechanism

IBM Cloud Object Storage Manager

Admin | Help | Sign Out

Monitor Configure Security Maintenance Administration

Enable / Disable Authentication Mechanisms

Configure whether users can access vault data using username/password authentication.

☒ Enable password authentication

Configure whether users can access vault data using access key authentication.

☒ Enable access key authentication

▲ Enabling "Hide secret access keys" will make all new or existing Secret Access Keys inaccessible on this page and all APIs. Secret Access Keys will only be visible once during creation. After this feature is turned on, it cannot be turned off unless all Access Keys are deleted in the system.

☐ Hide secret access keys

Cancel Update

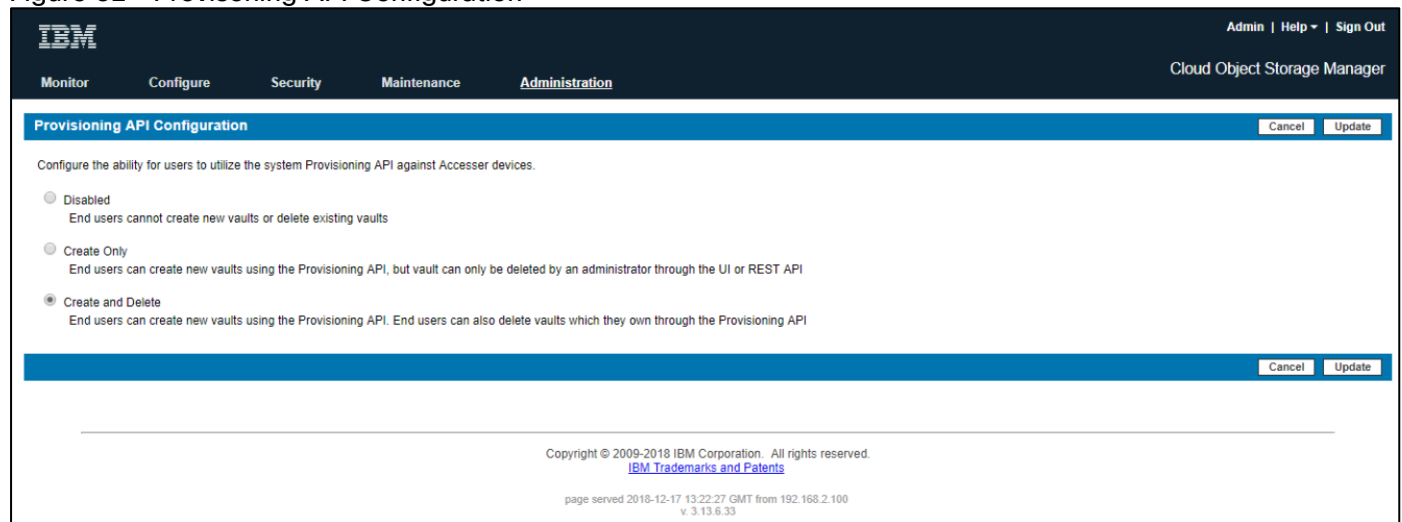
Copyright © 2009-2018 IBM Corporation. All rights reserved.
[IBM Trademarks and Patents](#)
 page served 2018-12-17 13:13:14 GMT from 192.168.2.100
 v. 3.13.6.33

Configure IBM COS Provisioning API

In order to create new vaults (sometimes referred to as buckets), it is important to enable the Provisioning API. In order to configure the IBM COS Provisioning API, follow these steps:

1. Open a web browser and navigate to the IP address of the IBM COS Manager 192.168.2.100 and log in with the credentials created previously.
2. Select the **Administration** tab from the topmost navigation bar and scroll down to Provisioning API Configuration and select Configure.
3. Select the bottom radio button next to Create and Delete. Next, select the Update button at the bottom right.

Figure 52 Provisioning API Configuration



Create a Storage Pool

A storage pool is defined by a logical grouping of Slicestor devices used to store vault data. A vault is initially created on a storage pool, and then may be expanded by creating new storage pool or pools on additional devices. Additional pools must be a multiple width of the original pool. A Slicestor device may only be a member of a single storage pool. To create a storage pool, follow these steps:

1. Select the **Configure** tab from the topmost navigation bar.
2. Select **Create Storage Pool** from underneath the **Summary** section.
3. Provide the Storage Pool a name in the **Name** field.
4. Select a width equal to the total number of Slicestors contained in the pool in the **width** drop-down box.
5. Selected the **Packed Storage** radio button which is ideal for objects that could be less than 32KB in size.
6. Select Enable the embedded Accesser service on all Slicestor devices belonging to this storage pool to enable Embedded Accesser.
7. Verify that all 3 Slicestors are selected in the **Devices** section and then select the **Save** button.

Figure 53 Create New Storage Pool

General

Name:

Width:

Storage Engine (cannot be changed later):

☒ **Packed Storage**
Recommended for all S3 or OpenStack use cases as well as any use case involving mirrors or small (<32 kB) objects.

☐ **File Storage**
Legacy engine for Simple Object use cases, particularly ones which include heavy delete activity.

☒ **Enable the embedded Accesser service on all Slicestor devices belonging to this storage pool**

Embedded Accesser Service Settings:

API Type:

API Ports:

- ☒ 80 HTTP
- ☒ 443 HTTPS
- ☒ 8080 HTTP
- ☒ 8443 HTTPS

Service API Ports:

- ☒ 8337 HTTP
- ☒ 8338 HTTPS

S3 Virtual Host Suffix: Example: *.ibm.com, *.ibm2.com

Additional Subject Alternative Names: Example: IP:10.0.0.1, IP:10.0.0.2, DNS:example1.domain.com

Suggest Devices

Devices

[Select All](#) [Unselect All](#)

Selected item count: 3

Name	Appliance Name	Drive Count	Site	Total Size	Version
<input checked="" type="checkbox"/> s slicestor1	Cisco UCS C240 M5	12	Cisco Valida...	119.06 TB	3.13.6.33
<input checked="" type="checkbox"/> s slicestor2	Cisco UCS C240 M5	12	Cisco Valida...	119.06 TB	3.13.6.33
<input checked="" type="checkbox"/> s slicestor3	Cisco UCS C240 M5	12	Cisco Valida...	119.06 TB	3.13.6.33

Create Vault for User Access

A Vault is a collection of data that is stored in one logical container, across a defined Storage Pool of Slicestor devices. Multiple Vaults may be linked to the same Storage Pool. There are several considerations for vault creation:

- In order to define a vault, the quantity of Slicestor devices (width) and the Threshold must be identified. The width and the threshold will interactively determine the maximum usable capacity. The number of devices in a Storage Pool must always be a multiple of the width, e.g. for a Storage Pool with 16 Slicestor devices, the width must be either 16 or 8.
- The vault threshold, always less than the width, will determine the reliability of the vault, i.e. how many slices must be minimally present to accurately read data. The Manager UI will allow any value between 1 and the Vault Width except for Dispersal Mode.
- The write threshold should be set larger than the threshold. If the number of slices available is less than or equal to the write threshold, the vault will be read-only. If the number of slices is greater than the write threshold but less than or equal to the alert threshold, the vault will remain fully functional but will trigger an alert.

To create a vault, follow these steps:

1. Select the **Configure** tab from the topmost navigation bar.
2. Select **Create Vault** from underneath the **Summary** section.

3. Provide a Vault name in the Name field.
4. Make certain to leave the check box next to **Enable Secure Slice Technology** checked.
5. Leave all other options at default and make any additional desired changes.
6. Select the Save button from the bottom right hand side.

Figure 54 Create New Standard Vault

IBM Admin | Help | Sign Out
Cloud Object Storage Manager

Monitor **Configure** Security Maintenance Administration

Storage

Name: Cisco-Pool
Width: 3
Devices: 3
Sites: Cisco Validated Design
Embedded Accesser Service: Enabled

Storage Capacity

0 bytes Used	357.18 TB Free	357.18 TB Raw Capacity
-----------------	-------------------	---------------------------

Create New Standard Vault Cancel Save

General

Name: * Cisco-Vault
Description: (optional)
Tags: Select one or more tags...

Configuration

A vault optimization is selected for you based on the storage pool width and device models.

Vault Optimization:

☒ Storage Efficiency ⓘ More usable capacity with reasonable performance. ☐ Performance ⓘ Better performance with less usable capacity.

Options

☒ Enable SecureSlice™ Technology
☐ Enable Versioning
☐ Delete Restricted
☐ Enable Server Side Encryption with Customer-Provided Keys (SSE-C)

Quotas

Soft Quota: (optional) TB
Hard Quota: (optional) TB

Advanced Index Settings

☒ Name Index Enabled
The index is needed to provide prefix-based listing and sorted listing results for named object vaults.
☐ Recovery Listing Enabled
Recovery listing allows for deterministic but unsorted listing results when the Name Index is disabled or corrupted. Some clients or application software may not function properly with unsorted listing results.

* Required field

Cancel Save

Create Vault Template for Provisioning

A Vault Template can be useful for vault or bucket creation via API. The first step for enabling this functionality is to create a vault.

To create a vault template, follow these steps:

1. Select the **Configure** tab from the topmost navigation bar.
2. Select **Configure** from underneath the **Template Management** section.
3. Within the Vault Template section beneath Template Management, select Cisco-Pool from the drop-down next to Select Storage Pool for Vault Template. Next, select **Create**.
4. Provide a Vault Template name in the **Name** field.

5. Make certain to leave the check box next to `Enable Secure Slice Technology` checked.
6. Select the checkbox next to the previously created Access Pool, `Cisco-Access`.
7. Leave all other options at default and make any additional desired changes.
8. Select the `Save` button from the bottom right hand side.
9. Return to the Template Management section by selecting the `Configure` tab from the topmost navigation bar and then selecting `Configure` from underneath the `Template Management`.
10. Select the radio button next to the newly created vault and then select `Update`.

Create an Access Pool

An Access Pool is a collection of Accessers that is user configurable to provide access to a vault or set of vaults. Users can be aware of which Accesser in the pool is providing connection, or a load balancer can be configured to automatically distribute load in a round robin fashion.

To create an Access Pool, follow these steps:

1. Select the `Configure` tab from the topmost navigation bar.
2. Select `Create Access Pool` from underneath the `Summary` section.
3. Provide an Access Pool name in the `Name` field.
4. From the API Type drop-down, select `Cloud Storage Object`.
5. Select `Cisco-Pool` from the `Storage Pool` field.
6. Select `Standard Vault` from the `Item Type` field.
7. Select the previously created `Cisco-Vault` at the bottom of the `Deployment` section.
8. Select the `Save` button from the bottom right side.

Figure 55 Create New Access Pool

IBM Cloud Object Storage Manager

Admin | Help | Sign Out

Monitor **Configure** Security Maintenance Administration

GO

Create New Access Pool Cancel Save

General

Name: Cisco-Access

Description:

API Type: Cloud Storage Object

API Ports:

- ☒ 80 HTTP
- ☒ 443 HTTPS
- ☒ 8080 HTTP
- ☒ 8443 HTTPS

Service API Ports:

- ☒ 8337 HTTP
- ☒ 8338 HTTPS

S3 Virtual Host Suffix: Example: *.ibm.com, *.ibm2.com

Additional Subject Alternative Names: Example: IP: 10.0.0.1, IP: 10.0.0.2, DNS: example1.domain.com

☒ Include default IPs in Subject Alternative Names

Access Devices

There are no Accesser devices available.

Deployment

Storage Pool: Cisco-Pool

Item Type: Standard Vault

Text Search: Search results... Clear filters

To select multiple items at once click on the desired check boxes while holding shift key.

Select All Unselect All

☒ Cisco-Vault

Visible items: 1 - Filtered from: 2
Selected item count: 1

Cancel Save

Screenshot

Copyright © 2009-2018 IBM Corporation. All rights reserved.
IBM Trademarks and Patents

Create a User for Object Access

It will be necessary to create an additional user for object access. There are protections in place to keep system administrators from being able to access data to avoid system compromise.

To create a new user, follow these steps:

1. Select the **Security** tab from the topmost navigation bar.
2. Select **Create Account** from underneath the **Accounts** section.
3. Provide a name in the **Name** field.
4. Provide a username in the **Username** field.
5. Provide a desired password in the **Password** and **Confirm Password** fields.
6. Select **Vault Provisioner** as Role.
7. Select the previously created **Cisco-Vault** at the bottom of the **Vault Access** section.
8. Select the button **Move to Owner**.

Figure 56 Create New Account

IBM Admin | Help | Sign Out
Cloud Object Storage Manager

Monitor Configure **Security** Maintenance Administration

Edit Account: Cisco-User Cancel Update

Name: Cisco-User
Email:

☒ Allow authentication with a username and password maintained within the Cloud Object Storage Manager
 Username: Cisco

Time Zone:
☒ Use the default manager time zone (GMT)
☐ Use a custom time zone for this account
 -- Select a time zone --

Roles

Assign Role	Read Only	Role	Description
<input type="checkbox"/>		Super User	Perform any action within the Cloud Object Storage Manager except vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	System Administrator	Perform any action within the Cloud Object Storage Manager except security, account management and vault read/write.
<input type="checkbox"/>	<input type="checkbox"/>	Security Officer	Perform security and account management actions within the Cloud Object Storage Manager.
<input type="checkbox"/>		Operator	Perform monitoring actions within the Cloud Object Storage Manager.
<input checked="" type="checkbox"/>		Vault Provisioner	Create / delete vaults using the Provisioning API. This role alone does not grant access to the Cloud Object Storage Manager interface.

Vault Access

*Move vaults between tabs to change their access permissions - then Update to save.

Owner (0) Read/Write (0) Read-Only (0) No Access (2) Show Filters

Select All Unselect All

Move to Owner Move to Read/Write Move to Read Only

☐ dsmgmt-cisco-pool
☒ Cisco-Vault

9. Select the Owner tab underneath Vault Access and verify that Cisco-Vault has been properly moved.

10. Select the Save button from the bottom right side.

Figure 57 Access for Owner

Vault Access

*Move vaults between tabs to change their access permissions - then Save.

Owner (1) Read/Write (0) Read-Only (0) No Access (1) Show Filters

Select All Unselect All

Move to Read/Write Move to Read Only Move to No Access

☒ Cisco-Vault

Cancel Save

11. Select the Security tab from the topmost navigation bar again if needed after creating the new user account.

12. Select the newly created user from beneath the Accounts section.

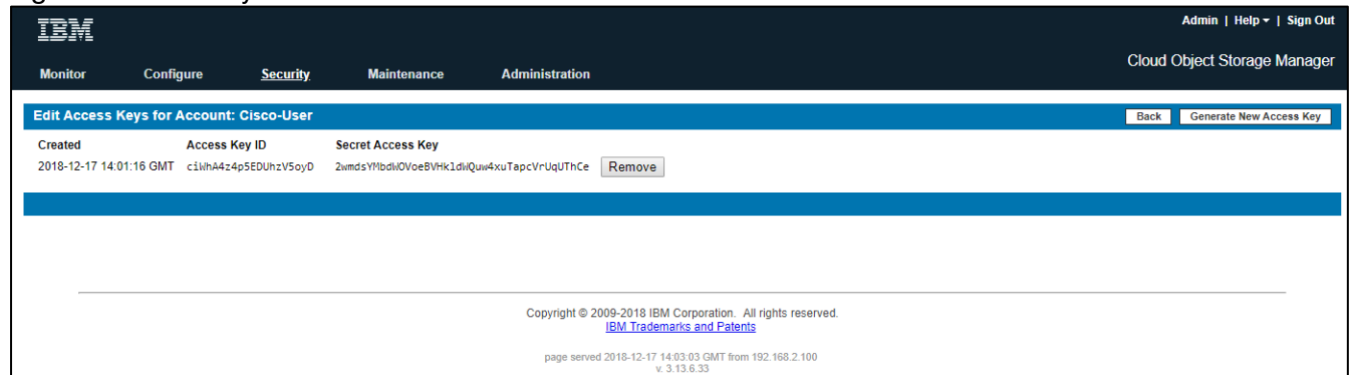
13. Find the Access Key Authentication section and select the Change Keys button.

14. From the Edit Access Keys for Account screen, select the Generate New Access Key button.

15. Once the key is created, select the Click to Show Secret Access Key button.

16. Make note of the Access Key ID and the Secret Access Key.

Figure 58 Edit Key Accessfor Account: Cisco-User



Functional Object Storage Access Validation

The next section describes one of the methods to access the newly created vault. For many end users, access may utilize the API, a GUI based program, or a program from the command line. For simplicity, a command line example is provided below.

Prerequisites

- A Linux system on the same network as the IBM COS Embedded Accessers (a virtual machine will suffice)
- The 'awscli' tool. Installation of this tool will be demonstrated with Ubuntu based 'apt'
- Access to either a local repository or the internet for remote application installation

To test IBM COS access functionality, follow these steps:

1. Update the apt repositories providing user password where necessary:

```
apt -y update
```

2. Install 'awscli'

```
apt -y install awscli
```

3. Create a new configuration file at ~/.aws/credentials with the following information. Use whichever tool is preferred, such as 'vim', however for simplicity, the file and credentials are created using the echo command. Modify access key id and secret access key as necessary:

```
echo [default] > ~/.aws/credentials
```

```
echo aws_access_key_id = ciWhA4z4p5EDUhzV5oyD >> ~/.aws/credentials
```

```
echo aws_secret_access_key = 2wmdsYmBdWVoeBVHkldWQuw4xuTapcVrUqUThCe >>
~/.aws/credentials
```

4. Test credentials, access, and the presence of the Cisco-Vault vault with the following command and output:

```
aws --endpoint-url=http://192.168.2.101 s3 ls
```

```
2018-12-16 13:11:21 Cisco-Vault
```

- List the contents of the Cisco-Vault vault (return output should be empty) :

```
aws --endpoint-url=http://192.168.2.101 s3 ls Cisco-Vault
```

- Upload an ISO file to the vault101 bucket

```
aws --endpoint-url=http://192.168.2.101 s3 cp /images/ubuntu-17.04-server-  
amd64.iso s3://Cisco-Vault/
```

```
upload: ../../images/ubuntu-17.04-server-amd64.iso to s3://Cisco-Vault/ubuntu-  
17.04-server-amd64.iso
```

- List vault contents and observe the presence of the newly uploaded ISO file:

```
aws --endpoint-url=http://192.168.2.101 s3 ls Cisco-Vault  
2018-12-16 15:41:35 718274560 ubuntu-17.04-server-amd64.iso
```

- Create a new vault or bucket:

```
aws --endpoint-url=http://192.168.2.101 s3api create-bucket --bucket Cisco-Vault1
```

- Copy the previously uploaded ISO file from Cisco-Vault bucket to Cisco-Vault1 bucket:

```
aws --endpoint-url=http://192.168.2.101 s3 cp s3://Cisco-Vault/ubuntu-17.04-  
server-amd64.iso s3://Cisco-Vault1/
```

```
copy: s3://Cisco-Vault/ubuntu-17.04-server-amd64.iso to s3://Cisco-Vault1/ubuntu-  
17.04-server-amd64.iso
```

- Move the newly copied ISO on Cisco-Vault1 back to Cisco-Vault under a new name:

```
aws --endpoint-url=http://192.168.2.101 s3 mv s3://Cisco-Vault1/ubuntu-17.04-  
server-amd64.iso s3://Cisco-Vault/ubuntu.iso
```

```
move: s3://Cisco-Vault1/ubuntu-17.04-server-amd64.iso to s3://Cisco-  
Vault/ubuntu.iso
```

- Verify that Cisco-Vault1 is now empty and the existence of two equally sized Ubuntu ISO files:

```
aws --endpoint-url=http://192.168.2.101 s3 ls s3://Cisco-Vault1
```

Initial Performance S3 Benchmark with COSBench

To verify the performance of the 3-node IBM COS cluster, a performance benchmark is initiated with COSBench¹ for write and read operations. Both tests are identically configured:

- Object size = 4 MB
- Clients = 3
- Workers per client = 32
- Ratio for read/write = 100%

The results are shown below.

¹ <https://github.com/intel-cloud/cosbench>

Table 8 Initial Performance Benchmark for S3 with COSBench

	Read	Write
Average Response Time	259 ms	573 ms
Average Bandwidth	1515 MB/s	710 MB/s

Validation

IBM COS High Availability Testing

It is important for business continuity to help ensure high availability of the hardware and software stack. Some of these features are built into the Cisco UCS Infrastructure and enabled by the software stack and some of these features are possible from the IBM COS Storage software itself. In order to properly test for high availability, the following considerations were given priority:

- The IBM COS deployment will process a reasonable amount of load when the fault is triggered. Total throughput will be recorded from the COSBench interface.
- Only a single fault will be triggered at any given time. Double failure is not a part of this consideration.
- Performance degradation is acceptable and even expected, but there should be no business interruption tolerated. The underlying infrastructure components should continue to operate within the remaining environment.
- The tests were conducted with the IBM COS configuration with Embedded Accesser.

The following High Availability tests were performed:

- Cisco Nexus 9332 Switch A failure
- Cisco UCS 6332 Fabric Interconnect A failure
- Cisco UCS C240 M5L – IBM COS slicestor1 disk failure
- Cisco UCS C240 M5L – IBM COS slicestor1 node failure

As indicated previously, a reasonable amount of load will be defined as follows:

- The COSBench application will be configured to send a steady stream of data to slicestor1.

Cisco Nexus C9332PQ High Availability Testing

Sequence of Events

1. Connect to Cisco Nexus 9332 Switch A and make certain running-config is copied to startup-config to make certain no configuration changes are lost during power cycle.

```
N9k-RH-A# copy run start
```

```
[#####] 100%
```

```
Copy complete.
```

2. Initiate load to the cluster by utilizing COSBench.
3. Pull out the power cables from Nexus switch N9k-RH-A and wait for at least 5 minutes before plugging in the power cables.



The load continues during Cisco Nexus C9332PQ reboot process.

Aside from loss of response from Nexus 9332 switch, IBM COS environment remained functional, load continued at constant rate, and redundancy was reestablished upon Switch A completing the reboot process.

Cisco UCS Fabric Interconnect 6332 High Availability Testing

Sequence of Events


1. Connect to Cisco UCS Manager and verify that both Fabric Interconnects and all nodes are in a known good working condition.
2. Initiate load to the cluster by utilizing COSBench.
3. Initiate reboot of Fabric Interconnect A. Establish a secure shell session to Fabric Interconnect A and enter the following commands.

```
connect local-mgmt
```

```
reboot
```

4. The Fabric Interconnect can take as long as 10 minutes to completely initialize after a reboot. Wait the entire amount of time for this process to complete.

The graph below is a snapshot from COSBench. At the vertical red line is where Fabric Interconnect A was rebooted. Only a slight loss in throughput was observed that could be within the noise of run-to-run variation. The total workload took place over the course of 15 minutes with ample time for Fabric Interconnect A to properly return to a known good state.




Cisco UCS C240 M5L Disk Failure Testing

Sequence of Events

1. Connect to Cisco UCS Manager and verify that both Fabric Interconnects and all nodes are in a known good working condition.
2. Initiate load to the cluster by utilizing COSBench.
3. Pull out one of 10 TB disks of slicestor1 and wait for at least 10 minutes.

The graph below is a snapshot from COSBench. At the vertical first red line is where Disk 1 was pulled. There was only a minimal drop at all and the overall write speed remained consistent. At the vertical second line is where disk was plugged in again. Overall there was only a very minimal drop in write performance.




Cisco UCS C240 M5L Node Failure Testing

Sequence of Events

1. Connect to Cisco UCS Manager and verify that both Fabric Interconnects and all nodes are in a known good working condition.
2. Initiate load to the cluster by utilizing COSBench.
3. Remove network cables from slicestor1 and wait for at least 10m minutes.

The graph below is a snapshot from COSBench. At the first vertical red line is where slicestor1 was rebooted. Only a slight loss in throughput of about 10 percent was observed. At the second red line the reboot of slicestor1 was finished and workload returned to a normal state. The total workload took place over the course of 15 minutes with ample time for slicestor1 to properly return to a known good state.



Summary

Object storage is an increasingly popular form of distributing data in a scale-out system. The entry size sinks to more and more smaller units. IBM with IBM COS is leading the pack with the new Concentrated Dispersal Mode technology when it comes to storing data as an object with high availability and reliability on small entry-level solutions.

The entry-level solution in this CVD provides customers and partners with everything necessary to store object data easily and securely. Cisco's leading technology of centralized management and advanced networking technology helps to easily deploy, manage and operate the IBM COS solution with Embedded Accessor in Concentrated Dispersal Mode.

Cisco and IBM are a perfect combination when it comes to reliably implementing new solutions in the area of object storage and showing customers new ways.

About the Authors

Oliver Walsdorf, Technical Marketing Engineer for Software Defined Storage, Computer Systems Product Group, Cisco Systems, Inc.

Oliver has more than 20 years of storage experience, working in different roles at different storage vendors, and is now an expert for software-defined storage at Cisco. For the past four years Oliver was focused on developing storage solutions at Cisco. He now works on IBM COS, develops Co-Solutions with IBM for the overall storage market and published several Cisco documents. With his focus on SDS he drives the overall attention in the market for new technologies. In his leisure time, Oliver enjoys hiking with his dog and motorcycling.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the following for their significant contribution and expertise that resulted in developing this document:

- Chris O'Brien, Cisco Systems, Inc.
- Jawwad Memon, Cisco Systems, Inc.
- Ulrich Kleidon, Cisco Systems, Inc.