



VersaStack Design Guide for IBM Cloud Object Storage with Cisco UCS S3260 Storage Server

Last Updated: October 30, 2018



About the Cisco Validated Design (CVD) Program

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

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

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

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

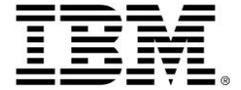
CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing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

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

© 2017 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	4
Solution Overview	6
Introduction	6
Solution	6
Audience	7
Solution Summary	7
Technology Overview	8
Cisco Unified Computing System	8
Cisco UCS Differentiators	9
Cisco UCS S3260 Storage Server	10
Cisco UCS C220 M4 Rack Server	11
Cisco UCS Virtual Interface Card 1387	12
Cisco UCS Fabric Interconnects	13
Cisco UCS 6300 Series Fabric Interconnect	14
Cisco Nexus 9332PQ Switch	15
Cisco UCS Manager	16
IBM Cloud Object Storage Architecture	17
Solution Design	19
Solution Architecture Overview	19
Compute Layer Design	20
Cisco UCS Server Connectivity to Unified Fabric	21
Validated Compute Design	21
High Availability	21
QoS and Jumbo Frames	21
Software Distributions and Versions	22
Hardware Requirements	22
Summary	26
About the Authors	27
Acknowledgements	27



Executive Summary

The data center of today continues to evolve to meet a variety of challenges that are no longer satisfied with traditional storage. Legacy architecture based on block and file storage face significant limitations that are not easily addressed without a radically different methods. Software Defined Storage addresses these limitations for a number of reasons:

- Software Defined Storage offers limitless scale and a decrease in management complexity.
- Software Defined Storage offers a simplified cost structure that is well suited for large-capacity needs. As the cost per gigabyte continues to shrink, Software Defined Storage becomes increasingly well suited for backup, archive, and cloud operations.
- Software Defined Storage for Object Storage breaks the monolithic storage mold. Object storage combines data, metadata, and unique identification to create objects.
- Software Defined Object Storage also provide better security and availability

Classic enterprise storage systems are designed to address business-critical requirements in the data center. And they still excel at this task today, but new trends and changing uses cases such as backup, active archive, and file-sync-and-share require new solutions built on new technology. Unstructured data aims to provide massive amounts of storage at extreme scale, particularly in environments where performance is less critical.

IBM Cloud Object Storage (IBM COS) is a software defined storage solution that brings massive scale and easy management to your datacenter. Reduced costs, tremendous scale, security, and enterprise grade reliability and availability is to be expected from this leading edge storage architecture.

The Cisco UCS S3260 Storage Server, designed **without compromise for today's** data center, combined with IBM COS dsNet is ideal for object storage solutions that require new demands in a world of unstructured data where data creation shows no signs of slowing - whether that workload is cloud data, a file repository, an active backup, or long term cold storage. The S3260 Storage Server provides a robust, comprehensive framework with unparalleled storage scalability combined with a standard 40 Gigabit Ethernet networking that allows Cisco to continue to push beyond was is possible in the data center. The S3260 Storage Server is the ideal object storage platform of choice specifically because of key differentiation from the competition:

- Proven industry standard server with a modular infrastructure and field upgradable components that reduce or eliminate the need for migration
- High-bandwidth networking that meets the needs of large-scale object storage solutions like IBM Cloud Object Storage
- Unified, embedded management for easy-to-scale infrastructure

Cisco and IBM are synergizing like never before to offer customers a scalable object storage solution for unstructured data that is integrated with IBM Cloud Object Storage. By leveraging the strength of Cisco UCS infrastructure and management, this solution is cost effective to deploy, easy to manage, and built on futureproof technology that will provide customers with the necessary tools for next generation cloud deployments.

Cisco Validated Designs (CVDs) include platforms and solutions that are designed, tested, and documented to improve customer deployments. These designs include a wide range of technologies and products into a portfolio of solutions that address the business needs of customers.

Solution Overview

Introduction

For years, traditional storage systems were able to easily service the demands of existing workloads. But in the past decade, much of the information housed on classic storage systems has been moved to an unstructured data architecture. In fact, almost eighty percent of data is currently unstructured today. This creates new opportunities to scale at a rate that entirely matches the consumption demand. Object storage is the most recent approach for managing these tremendous amounts of data.

IBM Cloud Object Storage provides organizations the flexibility, scalability and simplicity required to store, **manage and access today's rapidly growing unstructured data in a cloud environment.** Our time-tested solutions turn storage challenges into business advantage by reducing storage costs while reliably supporting both traditional and emerging cloud-born workloads for enterprise mobile, social, analytics and cognitive computing. IBM Cloud Object Storage is built on technology from object storage leader Cleversafe, acquired by IBM in 2015.

Scale-out object storage relies on traditional storage-optimized, x86 servers to reduce cost and increase performance. The Cisco UCS S3260 Storage Server is ideal for object storage solutions. It reduces deployment costs and leverages the power of the Cisco Unified Computing System (Cisco UCS). Management functions and features powered by Cisco UCS are unmatched by traditional unmanaged and agent-based systems. The S3260 Storage Server is ultimately the most flexible platform available that can be optimized for throughput, capacity, or compute intensive workloads.

By combining the massive scale that IBM Cloud Object Storage provides with the cost-effective simplicity of the Cisco S3260 Storage server, this solution delivers an enterprise scale-out storage architecture that is simple, fast, and completely scalable.

Solution

The current Cisco Validated Design (CVD) is a simple and linearly scalable architecture that provides an unstructured data solution on IBM Cloud Object Storage dsNet and Cisco UCS S3260 Storage Server. The solution includes the following features:

- Infrastructure for large scale object storage
- Design of a IBM Cloud Object Storage solution together with Cisco UCS S3260 Storage Server
- Simplified infrastructure management with Cisco UCS Manager
- Architectural scalability – linear scaling based on network, storage, and compute requirements
- Operational Guidance for properly sizing the IBM Cloud Object Storage architecture to fully leverage the benefits and features of Cisco UCS infrastructure

Audience

This document describes the architecture, design, and deployment procedures of an IBM Cloud Object Storage solution. The solution utilizes six Cisco UCS S3260 Storage Servers each configured with two M4 **Server Nodes, as IBM COS Slicestor's. In addition, depending on throughput needs, it leverages between** three to eight Cisco UCS C220 M4S rack-mount servers as IBM COS Accesser nodes and a single Cisco UCS C220 M4S rack-mount server as an IBM COS Manager node. All servers are managed by Cisco UCS Manager on two Cisco UCS 6332 Fabric Interconnects. 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 intend to deploy IBM Cloud Object Storage managed by the Cisco Unified Computing System (Cisco UCS) built upon the Cisco UCS S3260 Storage Server.

Solution Summary

This CVD describes in detail the process of deploying IBM Cloud Object Storage 3.10 on Cisco UCS S3260 Storage Server.

The configuration uses the following architecture for the deployment:

- 6 x Cisco UCS S3260 Storage Server with 2 x C3x60 M4 server nodes working as IBM COS Slicestor nodes
- 3-8 x Cisco UCS C220 M4S rack server working as IBM COS Accesser nodes
- 1 x Cisco UCS C220 M4S rack server working as IBM COS Manager node
- 2 x Cisco UCS 6332 Fabric Interconnect
- 1 x Cisco UCS Manager
- 2 x Cisco Nexus 9332PQ Switches

Technology Overview

Cisco Unified Computing System

Cisco Unified Computing System™ (Cisco UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 or 40 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, infrastructure platform where all resources are managed through a single, unified management domain.

The Cisco Unified Computing System consists of the following subsystems:

Compute - The compute piece of the system incorporates servers based on latest Intel's x86 processors. Servers are available in blade and rack form factors, managed by the same Cisco UCS Manager software.

Network - The integrated network fabric in the system provides low-latency, lossless, 10Gbps or 40Gbps Ethernet fabric. LANs, SANs, and high-performance computing networks which are separate networks today are consolidated within the fabric. 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 - Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.

Management - The system uniquely integrates all the system components, enabling the entire solution to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configuration and operations. Cisco UCS Manager helps in increasing the IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies; they help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

Cisco Unified Computing System has revolutionized the way servers are managed in a data center. This next section takes a detailed look at the unique differentiators in Cisco UCS and Cisco UCS Manager. An overview of the key sub-components leveraged in this architecture are also provided.

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 Differentiators

- Embedded Management – Servers in the system are managed by embedded software in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
- Unified Fabric – There is a single Ethernet cable to the FI from the server chassis (blade or rack) for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.
- Auto Discovery – By simply inserting a blade server in the chassis or connecting a rack server to the FI, discovery and inventory of compute resource occurs automatically without any intervention. Auto-discovery combined with unified fabric enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily without additional connections to the external LAN, SAN and management networks.
- Policy Based Resource Classification – Once a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy based resource classification of Cisco UCS Manager.
- Combined Rack and Blade Server Management – Cisco UCS Manager can manage Cisco UCS B-series blade servers and Cisco UCS C-series rack servers under the same Cisco UCS domain. Along with stateless computing, this feature makes compute resources truly agnostic to the hardware form factor.
- Model based Management Architecture – Cisco UCS Manager architecture and management database is model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, and Templates – The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Loose Referential Integrity – In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts from different domains, such as network, storage, security, server and virtualization the flexibility to work independently to accomplish a complex task.
- Policy Resolution – In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with **specific name is found in the hierarchy of the root organization, then special policy named “default” is**

searched. This policy resolution logic enables automation friendly management APIs and provides great flexibility to owners of different organizations.

- Service Profiles and Stateless Computing – A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in Multi-Tenancy Support – The combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and a service profiles based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- Virtualization Aware Network – Cisco VM-FEX technology makes the access network layer aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-**profiles defined by the network administrators'** team. VM-FEX also off-loads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
- Simplified QoS – Even though Fibre Channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco UCS S3260 Storage Server

The Cisco UCS® S3260 Storage Server (Figure 1) is a modular, high-density, high-availability dual node rack server well suited for service providers, enterprises, and industry-specific environments. It addresses the need for dense cost effective storage for the ever-growing data needs. Designed for a new class of cloud-scale applications, it is simple to deploy and excellent for big data applications, software-defined storage environments and other unstructured data repositories, media streaming, and content distribution.

Figure 1 Cisco UCS S3260 Storage Server



Extending the capability of the Cisco UCS C3000 portfolio, the Cisco UCS S3260 helps you achieve the highest levels of data availability. With dual-node capability that is based on the Intel® Xeon® processor E5-2600 v4 series, it features up to 600 TB of local storage in a compact 4-rack-unit (4RU) form factor. All hard-disk drives can be asymmetrically split between the dual-nodes and are individually hot-swappable. The drives can be built-in in an enterprise-class Redundant Array of Independent Disks (RAID) redundancy or be in a pass-through mode.

This high-density rack server comfortably fits in a standard 32-inch depth rack, such as the Cisco® R42610 Rack.

The Cisco UCS S3260 is deployed as a standalone server in both bare-metal or virtualized environments. Its modular architecture reduces total cost of ownership (TCO) by allowing you to upgrade individual components over time and as use cases evolve, without having to replace the entire system.

The Cisco UCS S3260 uses a modular server architecture that, using Cisco's blade technology expertise, allows you to upgrade the computing or network nodes in the system without the need to migrate data migration from one system to another. It delivers:

- Dual server nodes
- Up to 36 computing cores per server node
- Up to 60 drives mixing a large form factor (LFF) with up to 14 solid-state disk (SSD) drives plus 2 SSD SATA boot drives per server node
- Up to 512 GB of memory per server node (1 terabyte [TB] total)
- Support for 12-Gbps serial-attached SCSI (SAS) drives
- A system I/O Controller with Cisco VIC 1300 Series Embedded Chip supporting Dual-port 40Gbps
- High reliability, availability, and serviceability (RAS) features with tool-free server nodes, system I/O controller, easy-to-use latching lid, and hot-swappable and hot-pluggable components

Cisco UCS C220 M4 Rack Server

The Cisco UCS® C220 M4 Rack Server (Figure 2) is the most versatile, general-purpose enterprise infrastructure and application server in the industry. It is a high-density two-socket enterprise-class rack server that delivers industry-leading performance and efficiency for a wide range of enterprise workloads, including virtualization, collaboration, and bare-metal applications. The Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of the Cisco Unified Computing System™ (Cisco UCS) to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' total cost of ownership (TCO) and increase their business agility.

Figure 2 Cisco UCS C220 M4 Rack Server



The enterprise-class Cisco UCS C220 M4 server extends the capabilities of the Cisco UCS portfolio in a 1RU form factor. It incorporates the Intel® Xeon® processor E5-2600 v4 and v3 product family, next-generation

DDR4 memory, and 12-Gbps SAS throughput, delivering significant performance and efficiency gains. The Cisco UCS C220 M4 rack-mount server delivers outstanding levels of expandability and performance in a compact 1RU package:

- Up to 24 DDR4 DIMMs for improved performance and lower power consumption
- Up to 8 Small Form-Factor (SFF) drives or up to 4 Large Form-Factor (LFF) drives
- Support for 12-Gbps SAS Module RAID controller in a dedicated slot, leaving the remaining two PCIe Gen 3.0 slots available for other expansion cards
- A modular LAN-on-motherboard (mLOM) slot that can be used to install a Cisco UCS virtual interface card (VIC) or third-party network interface card (NIC) without consuming a PCIe slot
- Two embedded 1Gigabit Ethernet LAN-on-motherboard (LOM) ports

Cisco UCS Virtual Interface Card 1387

The Cisco UCS Virtual Interface Card (VIC) 1387 (Figure 3) is a Cisco innovation. It provides a policy-based, stateless, agile server infrastructure for your data center. This dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP) half-height PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter is designed exclusively for Cisco UCS C-Series and S3260 Rack Servers. The card supports 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). **It incorporates Cisco's next-generation converged network adapter (CNA) technology and offers a comprehensive feature set, providing investment protection for future feature software releases.** The card can present more than 256 PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Figure 3 Cisco UCS Virtual Interface Card 1387

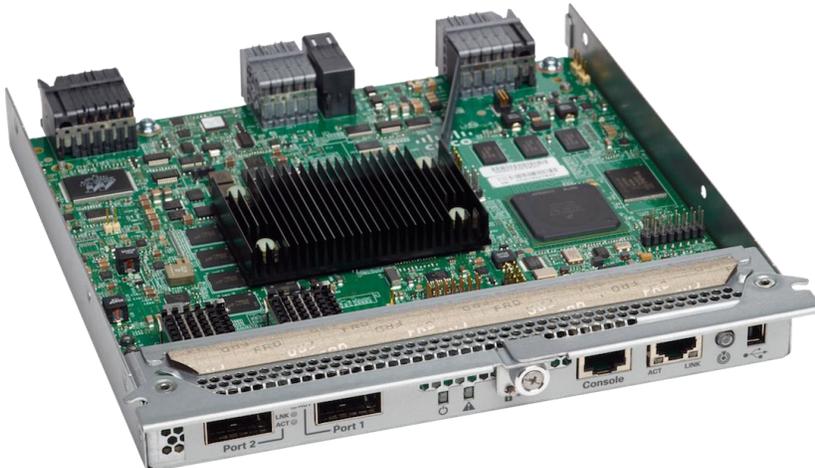


The Cisco UCS VIC 1387 provides the following features and benefits:

- Stateless and agile platform: The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure
- Network interface virtualization: Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect

The S3260 Storage Server uses a System I/O Controller (SIOC) with an embedded Cisco VIC 1300 series chip (Figure 4) that supports dual-port 40-Gbps connectivity. Both of the M4 Server Nodes within the S3260 Storage Server have a dedicated SIOC that supports all the same features that a standard mLOM or PCE-e 1300 Series VIC does. The three features of most significance to this design are the 40-Gbps connectivity, UCS Manager integration, and vNIC adapter profiles.

Figure 4 Cisco UCS SIOC with Embedded Virtual Interface Card



Cisco UCS Fabric Interconnects

The Cisco UCS Fabric interconnects provide a single point for connectivity and management for the entire system. Typically deployed as an active-active pair, **the system's fabric interconnects integrate all** components into a single, highly-available management domain controlled by Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in **deterministic I/O latency regardless of a server or virtual machine's topological location in the system.**

Fabric Interconnect provides both network connectivity and management capabilities for the Cisco UCS system. Cisco UCS Fabric Extenders (IOM) in the blade chassis support power supply, along with fan and blade management. They also support port channeling and, thus, better use of bandwidth. The IOMs support virtualization-aware networking in conjunction with the Fabric Interconnects and Cisco Virtual Interface Cards (VIC).

The capabilities of all Fabric Interconnects are summarized in Table 1 .

Table 1 Cisco UCS 6200 and 6300 series Fabric Interconnects

Features	6248	6296	6332	6332-16UP
Max 10G ports	48	96	96* + 2**	72* + 16
Max 40G ports	-	-	32	24
Max unified ports	48	96	-	16
Max FC ports	48 x 2/4/8G FC	96 x 2/4/8G FC	-	16 x 4/8/16G FC

* Using 40G to 4x10G breakout cables

** Requires QSA module

Cisco UCS 6300 Series Fabric Interconnect

The Cisco UCS 6300 Series Fabric Interconnects (Figure 5) are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 10 and 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

Figure 5 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, Cisco UCS 5100 Series Blade Server Chassis, and Cisco UCS C-Series Rack Servers managed by Cisco UCS. All servers attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all servers within its domain.

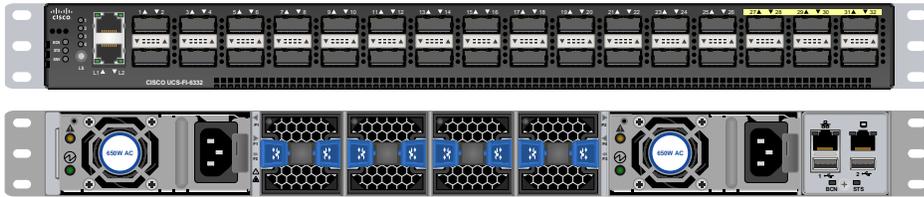
From a networking perspective, the Cisco UCS 6300 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 and 40 Gigabit Ethernet ports, switching capacity of 2.56 terabits per second (Tbps), and 320 Gbps of bandwidth per chassis, independent of packet size and enabled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

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 (Figure 6) 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.

Figure 6 Cisco UCS 6332 Fabric Interconnect – Front and Rear



Cisco Nexus 9332PQ Switch

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

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 (Figure 7) 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.

Figure 7 Cisco 9332PQ



All the 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.

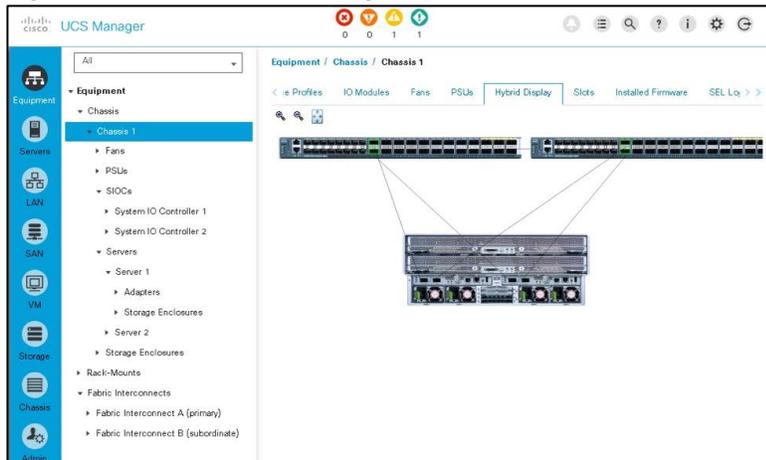
With the Cisco Nexus 9000 Series, organizations can quickly and easily upgrade existing data centers to carry 40 Gigabit Ethernet to the aggregation layer or to the spine (in a leaf-and-spine configuration) through advanced and cost-effective optics that enable the use of existing 10 Gigabit Ethernet 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 (Figure 8) provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) across multiple chassis, rack servers and thousands of virtual machines. It supports all Cisco UCS product models, including Cisco UCS B-Series Blade Servers, C-Series Rack Servers, 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 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 8 Cisco UCS Manager



An instance of Cisco UCS Manager with all Cisco UCS components managed by it forms a Cisco UCS domain, which can include up to 160 servers. In addition to provisioning Cisco UCS resources, this infrastructure management software provides a model-based foundation for 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. Cisco Manager provides flexible role- and policy-based management using service profiles and templates.

Service profiles benefit both virtualized and non-virtualized environments and increase the mobility of non-virtualized servers, such as when moving workloads from server to server or taking a server offline for service or upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.

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).

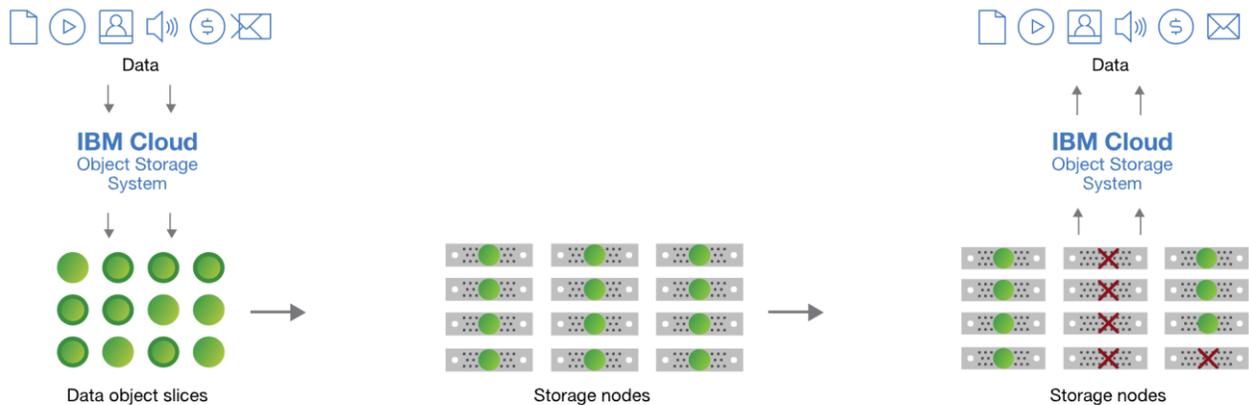
The Cisco UCS XML API provides comprehensive access to all Cisco UCS Manager functions. The API provides Cisco UCS system visibility to higher-level systems management tools from independent software vendors (ISVs) such as VMware, Microsoft, and Splunk as well as tools from BMC, CA, HP, IBM, and others.

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 Architecture

IBM Cloud Object Storage (Figure 9) is a dispersed storage mechanism that uses a cluster of storage nodes to store pieces of the data across the available nodes. IBM Cloud Object Storage uses an Information Dispersal Algorithm (IDA) to break files into unrecognizable slices that are then distributed to the storage nodes. No single node has all the data, which makes it secure and resilient to system failures. Security is made possible through encrypting data at rest to provide government grade security with built in key management. Additionally, no single disk, node or site contains enough information to constitute a data breach, while needing only a subset of the storage nodes to be available to fully retrieve the stored data. This ability to reassemble all the data from a subset of the chunks dramatically increases the tolerance to node and disk failures.

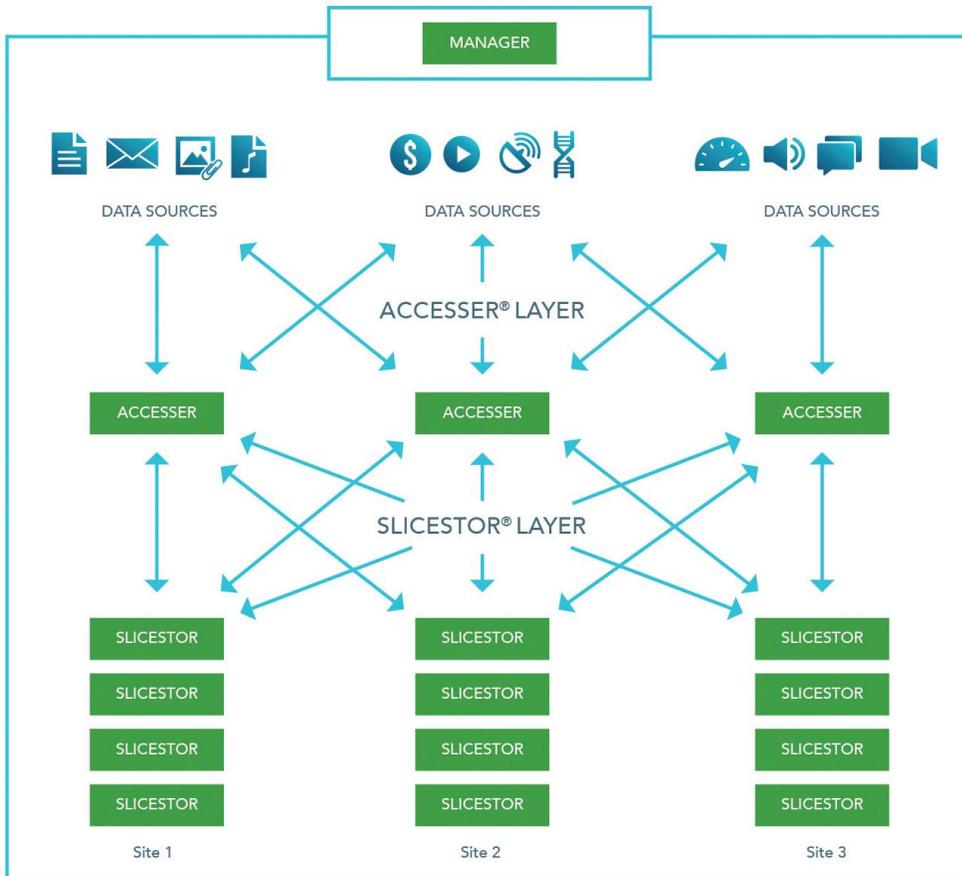
Figure 9 IBM Cloud Object Storage



The IBM Cloud Object Storage architecture is composed of three functional components. Each of these components runs ClevOS software that can be deployed on Cisco designed, scale-out storage optimized hardware. The three components include:

- IBM Cloud Object Storage Manager provides an out of band management interface that is used for administrative tasks, such as system configuration, storage provisioning, and monitoring the health and performance of the system
- IBM Cloud Object Storage Accesser imports and reads data, encrypting/encoding data on import and decrypting/decoding data on read. It is a stateless component that presents the storage interfaces to the client applications and transforms data by using an IDA
- The IBM Cloud Object Storage Slicestor node is primarily responsible for storage of the data slices. It receives data from the Accesser on import and returns data to the Accesser as required by reads

Figure 10 IBM dsNet Overview



- 1 dsNet[®] Manager Is deployed to configure and manage the infrastructure
- 2 Accessers are deployed to access the underlying storage
- 3 Slicestors are deployed to store data

Solution Design

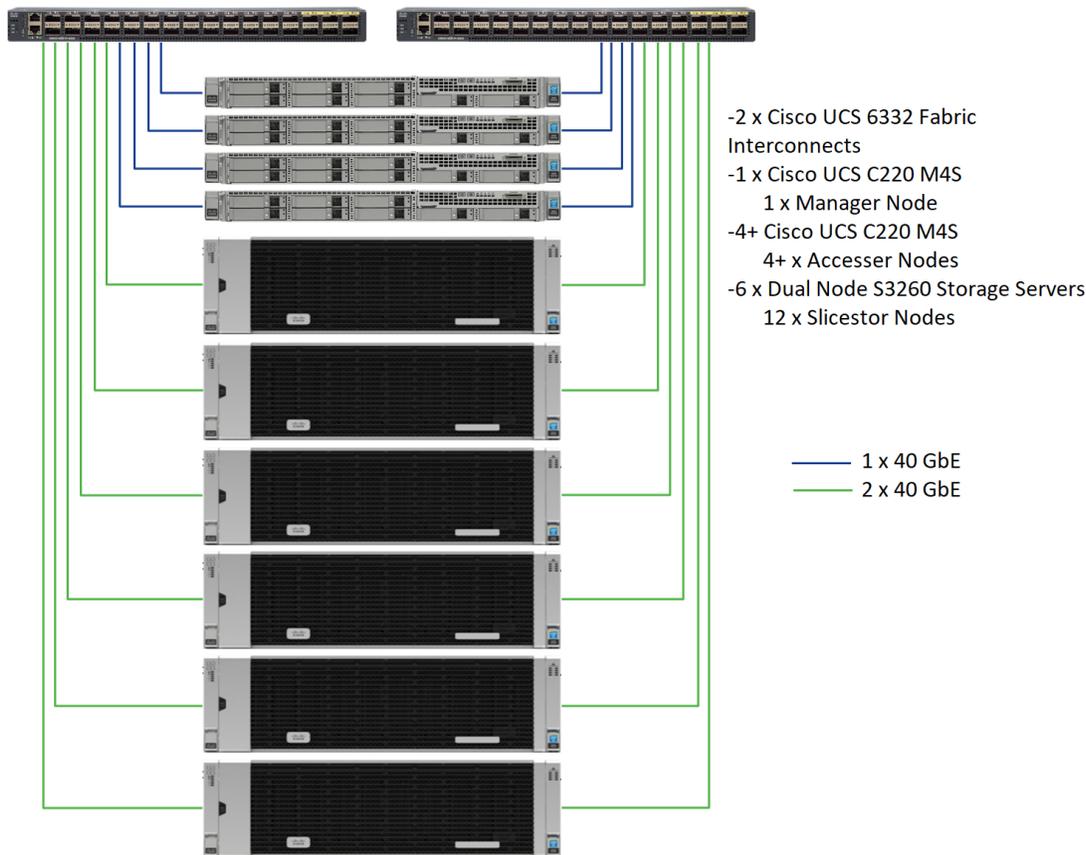
Solution Architecture Overview

This Cisco Validated Design provides a comprehensive, end-to-end guide for deploying IBM Cloud Object Storage on Cisco UCS S3260 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. All IBM Cloud Object Storage components – Manager, Accesser, and Slicestor – utilize the robust throughput and low latency only provided by the Cisco UCS 6332 Fabric Interconnect. Additionally, all components take advantage of the flexibility provided by Cisco UCS service profiles and service profile templates. By design, all design decisions and features within all in-use service profiles are able to be updated with a few modifications.

This design optionally 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).

Figure 11 Cisco UCS Software Defined Storage Architecture



Compute Layer Design

The compute resources supported in this design are Cisco UCS S3260 Storage Servers with accompanying M4 Server Nodes and Cisco UCS C220 M4S rackmount servers. Each Cisco UCS server is equipped with a Cisco Virtual Interface Card (VIC) that aggregate all traffic to and from the server across a single 40GbE interface. Cisco VICs eliminate the need for separate physical interface cards on each server for data and management connectivity. Cisco VIC adapter profiles are easily manageable from within the Cisco UCS Manager UI. Capabilities such as receive and transmit queues, receive side scaling, and jumbo frames are all configurable from within the management interface. This design guide focuses on a fairly basic, flat network topology, but some deployments might require a far more advanced configuration that isolates certain types of traffic. This is all easily achievable with the Cisco VIC without the purchase or configuration of any additional hardware. The Cisco VIC found in each server in this design guide is equipped with two 40 Gbps interfaces that provide best in class latency, throughput, and manageability. Cisco VICs can be virtualized to create up to 256 virtual interfaces that can be dynamically configured as virtual network interface cards (vNICs) or virtual host bus adapters (vHBAs). These virtual interfaces will appear as a standards-compliant PCIe endpoints to the OS. The scope of this solution with IBM Cloud Object Storage ClevOS is configured with two virtual NICs, one on each VIC interface. IBM COS is configured to leverage these two vNICs to provide operational active-backup redundancy in software.

Multiple models of Cisco VICs are available. Cisco VICs are available in the S3260 Storage Server as a part of the SIOC present on every model. The Cisco C220 M4S leverages a VIC that is available as a modular

LAN on Motherboard (mLOM). Additionally, a half-height PCI Express (PCIe) card is also available as an alternative configuration or for additional throughput.

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 the Cisco Unified Computing System, providing unified management and connectivity to all attached blades, chassis and rack servers. Fabric Interconnects provide a lossless and deterministic Fibre Channel over Ethernet (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

For validation, Cisco UCS S3260 Storage Servers with System IO Controller with included VIC 1380 and Cisco UCS C220 M4S servers with VIC 1387, were connected to 2 x Cisco UCS 6332 Fabric Interconnects. Each Cisco UCS C220 M4S was deployed with two 40 GbE QSFP copper cables. Each Cisco UCS M4 Server Node within the S3260 Storage Server was also deployed with two 40 GbE QSFP cables. Two M4 Server Nodes were within each S3260 Storage Server which results in a total capable throughput of 160 Gbps to each S3260 chassis.

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 system provides redundancy at the component and link level and end-to-end path redundancy to the LAN network.
- Cisco UCS S3260 storage server platform is highly redundant with redundant power supplies, fans and SIOC modules.
- 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 Manager, Accesser, and Slicestor 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.
- Each Manager, Accesser, and Slicestor is configured in mode 1 active-backup bonding mode at the ClevOS software layer.

QoS and Jumbo Frames

Cisco UCS, Cisco Nexus, and IBM Cloud Object Storage 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 2 .

Table 2 Software Versions

Layer	Component	Version or Release
Storage (Chassis) UCS S3260	Chassis Management Controller	3.0(3a)
	Shared Adapter	4.1(3a)
Compute (Server Nodes) UCS C3X60 M4	BIOS	C3X60M4.3.0.3b
	CIMC Controller	3.0(3a)
Compute (Rack Server) C220 M4S	BIOS	C220M4.3.0.3a
	CIMC Controller	3.0(3a)
Network 6332 Fabric Interconnect	UCS Manager	3.1(3a)
	Kernel	5.0(3)N2(3.13a)
	System	5.0(3)N2(3.13a)
Network Nexus 9332PQ	BIOS	07.59
	NXOS	7.0(3)I5(1)
Software	IBM COS ClevOS	3.10.0.126-ucs3

Hardware Requirements

This Cisco Validated Design describes the architecture, design, and deployment of an IBM Cloud Object Storage solution on six Cisco UCS S3260 Storage servers with two M4 Server Nodes, each configured as IBM COS Slicestor Nodes, and at least four Cisco UCS C220 M4S Rack Servers configured Accesser Nodes, and a single Cisco UCS C220 M4S Rack Server configured as a Manager node. The entire solution is connected to a pair of Cisco UCS 6332 Fabric Interconnects and a pair of Cisco Nexus 9332PQ switches with 40Gbps end-to-end.

The detailed configuration is as follows:

- Two Cisco Nexus 9332PQ Switches
- Two Cisco UCS 6332 Fabric Interconnects
- Six Cisco UCS S3260 Storage Servers with two M4 server nodes each

- Five Cisco UCS C220 M4S Rack Servers

Figure 12 Hardware Overview

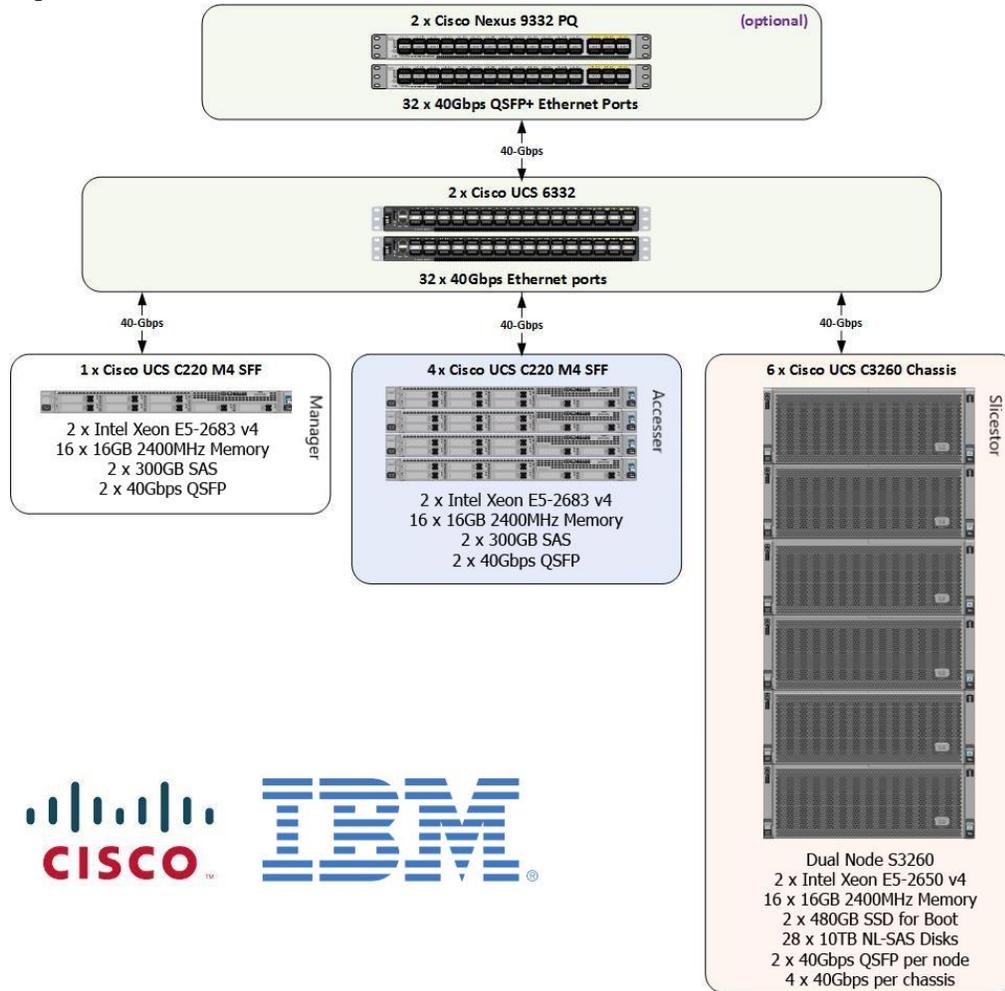


Table 3 List of Components

Component	Model	Quantity	Comments
IBM COS Slicestor	Cisco UCS S3260 M4 Chassis	6	<ul style="list-style-type: none"> 2 x UCS C3X60 M4 Server Nodes per Chassis (Total = 12 nodes) Per Server Node <ul style="list-style-type: none"> 2 x Intel E5-2650 v4, 256 GB RAM Cisco 12G SAS RAID Controller 2 x 480 GB SSD for OS, 26 x 10TB HDDs for Data

Component	Model	Quantity	Comments
			<ul style="list-style-type: none"> • Dual-port 40 Gbps VIC
IBM COS Ac- cesser	Cisco UCS C220M4S Rack server	4	<ul style="list-style-type: none"> • 2 x Intel E5-2683v4, 256 GB RAM • Cisco 12G SAS RAID Controller • 2 x 300 GB SAS for OS • Dual-port 40 Gbps VIC
IBM COS Man- ager	Cisco UCS C220M4S Rack server	1	<ul style="list-style-type: none"> • 2 x Intel E5-2683v4, 256 GB RAM • Cisco 12G SAS RAID Controller • 2 x 300 GB SAS for OS • Dual-port 40 Gbps VIC
Cisco UCS Fabric Inter- connects	Cisco UCS 6332 Fabric Interconnects	2	
Switches	Cisco Nexus 9332PQ Switches	2	

Summary

Cisco UCS S3260 Storage server is an ideal candidate for all IBM Cloud Object Storage deployments. The Cisco UCS 6332 Fabric Interconnect is an optimal infrastructure foundation to deploy IBM COS for maximum efficiency. Built with the latest generation of processors from Intel and years of Cisco DNA in the Virtual Interface Card, this solution is the most robust, agile, and manageable solution for scale out object storage. Together, IBM and Cisco have created a platform that is both flexible and scalable for multiple object storage use cases and applications. File sync and share, active-archive, or media streaming and collaboration, Cisco Unified Computing System and IBM Cloud Object Storage enables customers to right-size their infrastructure and adapt to their evolving business requirements.

About the Authors

Travis Hindley, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems, Inc.

Travis has almost 20 years of experience in systems engineering focusing on virtualization performance, server optimization, and storage solutions. As a member of the server performance engineering team on two **of the largest companies in today's datacenter**, Travis works to provide collateral that spans business units, partners, and competitors to bring the most exciting solutions to market and ultimately ensure customer success with technology in the datacenter solutions space.

Acknowledgements

- Ulrich Kleidon, Cisco Systems, Inc.
- Jawwad Memon, Cisco Systems, Inc.
- J.T. Wood, IBM
- Dan Albright, IBM