

Cisco UCS C240 M5 Server with Clouddian HyperStore Object Storage Design Guide

Last Updated: January 9, 2020



About the Cisco Validated Design Program

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

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

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

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

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 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)

© 2020 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction.....	6
Audience	6
Purpose of this Document.....	6
Solution Summary	6
Technology Overview	8
Cisco Unified Computing System	8
Cisco UCS Manager.....	8
Cisco UCS 6300 Fabric Interconnects.....	9
Cisco UCS C9336C-FX2 Nexus Switches	10
Cisco UCS C240 M5 Rack Server	11
Cisco UCS C220 M5 Rack-Mount Server	13
Cisco UCS Virtual Interface Card 1385.....	14
Red Hat Enterprise Linux 7.6.....	15
Cloudian HyperStore.....	16
Cloudian Object Storage.....	16
Cloudian HyperStore Design	17
Cloudian HyperStore Architecture	18
Cloudian Management Console.....	19
S3 Compatible.....	20
Integrated Billing, Management, and Monitoring	21
Infinitely Scalable on Demand.....	21
Security	21
Data Protection.....	21
Effortless Data Movement	21
Solution Design	22
Deployment Architecture	22
System Hardware and Software Specifications.....	23
Software Versions	23
Hardware Requirements and Bill of Materials	23
Physical Topology and Configuration	24
Network Topology	26
High Availability	27
Design Criteria	29

Requirements	29
Design Considerations.....	30
Number of Nodes of Cisco UCS C240 M5.....	30
Replication versus Erasure Coding	30
Flash Storage	33
JBOD versus RAID0 Disks	33
Memory Sizing	33
Network Considerations.....	33
Uplinks	33
Multi-Site Deployments.....	33
Expansion of the Cluster.....	34
Performance Characterization.....	36
Summary	40
About the Authors.....	41
Acknowledgements	41



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.

Most of the modern data centers are moving away from traditional file system type storage, to object storages. Object storage offers simple management, unlimited scalability and custom metadata for objects. With its low cost per gigabyte of storage, Object storage systems are suited for archive, backup, Life sciences, video surveillance, healthcare, multimedia, message and machine data, and so on.

Cisco and Cloudian are collaborating to offer customers a scalable object storage solution for unstructured data that integrates Cisco Unified Computing System (Cisco UCS) with Cloudian HyperStore. With the power of the Cisco UCS management framework, the solution is cost effective to deploy and manage and will enable the next-generation cloud deployments that drive business agility, lower operational costs and avoid vendor lock-in.

This validated design provides the framework for designing Cloudian HyperStore 7.1.4 on Cisco UCS C240 M5L Rack Servers. Cisco Unified Computing System provides the compute, network, and storage access components for the Cloudian HyperStore, deployed as a single cohesive system. The reference architecture described in this document is a realistic use case for deploying Cloudian HyperStore object storage on Cisco UCS C240 M5L Rack Server.

Solution Overview

Introduction

Object storage is a highly scalable system for organizing and storing data objects. Object storage does not use a file system structure, instead it ingests data as objects with unique keys into a flat directory structure and the metadata is stored with the objects instead of hierarchical journal or tree. Search and retrieval are performed using RESTful API's, which uses HTTP verbs such as GETs and PUTs. Most of the newly generated data, about 60 to 80 percent, is unstructured today and new approaches using x86 servers are proving to be more cost effective, providing storage that can be expanded as easily as data grows. Scale-out Object storage is the newest cost-effective approach for handling large amounts of data in the Petabyte and Exabyte range.

The Cloudian HyperStore is an enterprise object storage solution that offers S3 API based storage. The solution is highly scalable and durable. The software is designed to create unbounded scale-out storage systems that accommodates Petabyte scale data from multiple applications and use-cases, including both object and file-based applications. Cloudian Hyperstore can deliver a fully enterprise-ready solution that can manage different workloads and still remain flexible.

The Cisco UCS® C240 M5 Rack Server delivers industry-leading performance and expandability. The Cisco UCS C240 M5 rack server is capable of addressing a wide range of enterprise workloads, including data-intensive applications such as Cloudian HyperStore. The C240 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. The Cisco UCS brings the power and automation of unified computing to the enterprise, it is an ideal platform to address capacity-optimized and performance-optimized workloads.

This document describes the architecture and design procedures of Cloudian Hyperstore software on Cisco UCS C-Series Rack Servers.

Audience

The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineers, IT architects, and customers who want to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation. The reader of this document is expected to have the necessary training and background to install and configure Red Hat Enterprise Linux, Cisco Unified Computing System, Cisco Nexus and Cisco UCS Manager, as well as a high-level understanding of Cloudian Hyperstore Software and its components. External references are provided where applicable and it is recommended that the reader be familiar with these documents.

Readers are also expected to be familiar with the infrastructure, network and security policies of the customer installation.

Purpose of this Document

This document describes the steps required to design Cloudian HyperStore 7.1.4 on Cisco UCS platform. It discusses design choices and best practices using this shared infrastructure platform.

Solution Summary

Cisco and Cloudian developed a solution that meets the challenges of scale-out storage. This solution uses Cloudian HyperStore Object Storage software with Cisco UCS C-Series Rack Servers powered by Intel Xeon

processors. The advantages of Cisco UCS and Clouidian HyperStore combine to deliver an object storage solution that is simple to install, scalable, high performance, robust availability, system management, monitoring capabilities and reporting.

The configuration uses the following components for the deployment:

- Cisco Unified Computing System
 - Cisco UCS 6332 Series Fabric Interconnects
 - Cisco UCS C240 M5L Rack Servers
 - Cisco UCS Virtual Interface Card (VIC) 1385
 - Cisco C220M5 servers with VIC 1387
- Cisco Nexus 9000 C9336C-FX2 Series Switches
- Clouidian HyperStore 7.1.4
- Red Hat Enterprise Linux 7.6

Technology Overview

Cisco Unified Computing System

Cisco Unified Computing System 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 Processor scalable family. The Cisco UCS servers offer the patented Cisco Extended Memory Technology to support applications with large datasets and allow more virtual machines per server.
- Network - The system is integrated onto a low-latency, lossless, 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, 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.

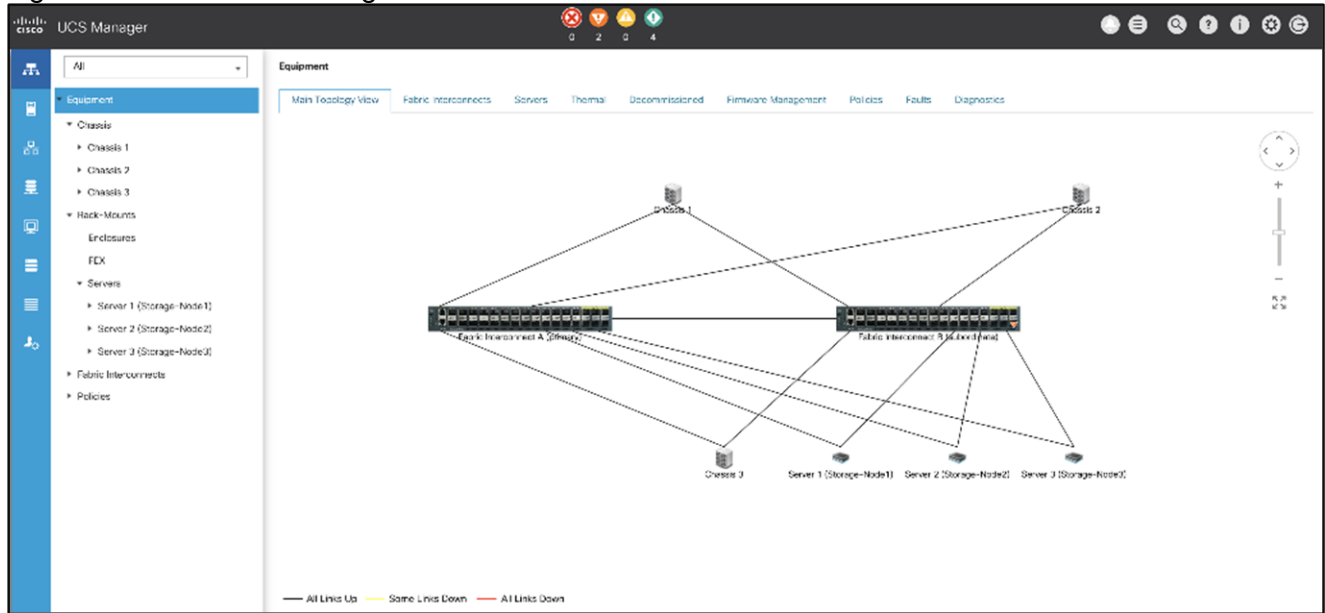
Cisco Unified Computing System is designed to deliver:

- 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 Manager

Cisco UCS Manager (UCSM) provides a unified, embedded management of all software and hardware components of the 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-Mount Servers, and Cisco UCS Mini, as well as the associated storage resources and networks. Cisco UCS Manager is embedded on a pair of Cisco UCS 6400, 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 1 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. 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 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). It can be integrated with Cisco Intersight which provides intelligent cloud-powered infrastructure management to securely deploy and manage infrastructure either as Software as a Service (SaaS) on Intersight.com or running on-premises with the Cisco Intersight virtual appliance.

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.

Cisco UCS 6300 Fabric Interconnects

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

Figure 2 Cisco UCS 6300 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, and C-Series Rack Servers managed by Cisco UCS. All servers attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all servers within its domain.

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

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 UCS C9336C-FX2 Nexus Switches

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.

Figure 3 Cisco Nexus C9336C-FX2 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 C9336C-FX2 Switch is a 1-rack-unit (1RU) switch that supports 7.2 Tbps of bandwidth and over 2.8 billion packets per second (bps) across thirty-six 10/25/40/100 -Gbps Enhanced QSFP28 ports

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 C240 M5 Rack Server

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

In response to ever-increasing computing and data-intensive real-time workloads, the enterprise-class Cisco UCS C240 M5 server extends the capabilities of the Cisco UCS portfolio in a 2RU form factor. It incorporates the Intel® Xeon® Scalable processors, supporting up to 20 percent more cores per socket, twice the memory capacity, and five times more.

Non-Volatile Memory Express (NVMe) PCI Express (PCIe) Solid-State Disks (SSDs) compared to the previous generation of servers. These improvements deliver significant performance and efficiency gains that will improve application performance. The C240 M5 delivers outstanding levels of storage expandability with exceptional performance, with:

- The latest second-generation Intel Xeon Scalable CPUs, with up to 28 cores per socket
- Supports the first-generation Intel Xeon Scalable CPU, with up to 28 cores per socket
- Support for the Intel Optane DC Persistent Memory (128G, 256G, 512G)[1]
- Up to 24 DDR4 DIMMs for improved performance including higher density DDR4 DIMMs
- Up to 26 hot-swappable Small-Form-Factor (SFF) 2.5-inch drives, including 2 rear hot-swappable SFF drives (up to 10 support NVMe PCIe SSDs on the NVMe-optimized chassis version), or 12 Large-Form-Factor (LFF) 3.5-inch drives plus 2 rear hot-swappable SFF drives
- Support for 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe Generation 3.0 slots available for other expansion cards
- Modular LAN-On-Motherboard (mLOM) slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot, supporting dual 10- or 40-Gbps network connectivity
- Dual embedded Intel x550 10GBASE-T LAN-On-Motherboard (LOM) ports
- Modular M.2 or Secure Digital (SD) cards that can be used for boot

Figure 4 Cisco UCS C240 M5L Front



Figure 5 Cisco UCS C240 M5 Internals



The Cisco UCS C240 M5 rack server is well-suited for a wide range of enterprise workloads, including:

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

Cisco UCS C240 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. When used in combination with Cisco UCS Manager, the Cisco UCS C240 M5 brings the power and automation of unified computing to enterprise applications, including Cisco® SingleConnect technology, drastically reducing switching and cabling requirements.

Cisco UCS Manager uses service profiles, templates, and policy-based management to enable rapid deployment and help ensure deployment consistency. It also enables end-to-end server visibility, management, and control in both virtualized and bare-metal environments.

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. The Cisco UCS C240 M5 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 C220 M5 Rack-Mount Server

The Cisco UCS C220 M5 Rack-Mount Server is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density 2-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. The Cisco UCS C-Series Rack-Mount Servers can be deployed as standalone servers or as part of Cisco UCS to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

The Cisco UCS C220 M5 server extends the capabilities of the Cisco UCS portfolio in a 1-Rack-Unit (1RU) form factor. It incorporates the Intel® Xeon Scalable processors, supporting up to 20 percent more cores per socket, twice the memory capacity, 20 percent greater storage density, and five times more PCIe NVMe Solid-State Disks (SSDs) compared to the previous generation of servers. These improvements deliver significant performance and efficiency gains that will improve application performance.

Figure 6 Cisco UCS C220M5 Rack-Mount Server



The Cisco UCS C220 M5 SFF server extends the capabilities of the Cisco Unified Computing System portfolio in a 1U form factor with the addition of the Intel Xeon Processor Scalable Family, 24 DIMM slots for 2666MHz DIMMs and capacity points up to 128GB, two 2 PCI Express (PCIe) 3.0 slots, and up to 10 SAS/SATA hard disk drives (HDDs) or solid state drives (SSDs). The Cisco UCS C220 M5 SFF server also includes one dedicated internal slot for a 12G SAS storage controller card.

The Cisco UCS C220 M5 server included one dedicated internal modular LAN on motherboard (mLOM) slot for installation of a Cisco Virtual Interface Card (VIC) or third-party network interface card (NIC), without consuming a PCI slot, in addition to 2 x 10Gbase-T Intel x550 embedded (on the motherboard) LOM ports.

The Cisco UCS C220 M5 server can be used standalone, or as part of the Cisco Unified Computing System, which unifies computing, networking, management, virtualization, and storage access into a single integrated architecture enabling end-to-end server visibility, management, and control in both bare metal and virtualized environments.

Cisco UCS Virtual Interface Card 1385

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

Figure 7 Cisco UCS VIC 1385



The Cisco UCS VIC 1385 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.

VIC 1385 has a hardware classification engine. This provides support for advanced data center requirements including stateless network offloads for NVGRE and VXLAN (VMware only), low-latency features for usNIC and RDMA, and performance optimization applications such as VMQ, DPDK, and Cisco NetFlow. The Cisco UCS VIC 1385 provides high network performance and low latency for the most demanding applications:

- Big data, high-performance computing (HPC), and high-performance trading (HPT)

- Large-scale virtual machine deployments
- High-bandwidth storage targets and archives

When the VIC 1385 is used in combination with Cisco Nexus® 3000 Series Switches, big data and financial trading applications benefit from high bandwidth and low latency. When the VIC is connected to Cisco Nexus 5000 Series Switches, pools of virtual hosts scale with greater speed and agility. The Cisco Nexus 6004 Switch provides native 40-Gbps FCoE connectivity from the VIC to both Ethernet and Fibre Channel targets.

Red Hat Enterprise Linux 7.6

Red Hat® Enterprise Linux is a high-performing operating system that has delivered outstanding value to IT environments for more than a decade. More than 90 percent of Fortune Global 500 companies use Red Hat products and solutions including Red Hat Enterprise Linux. As the world's most trusted IT platform, Red Hat Enterprise Linux has been deployed in mission-critical applications at global stock exchanges, financial institutions, leading telcos, and animation studios. It also powers the websites of some of the most recognizable global retail brands.

Red Hat Enterprise Linux:

- Delivers high-performance, reliability, and security
- Is certified by the leading hardware and software vendors
- Scales from workstations, to servers, to mainframes
- Provides a consistent application environment across physical, virtual, and cloud deployments

Designed to help organizations make a seamless transition to emerging datacenter models that include virtualization and cloud computing, Red Hat Enterprise Linux includes support for major hardware architectures, hypervisors, and cloud providers, making deployments across physical and different virtual environments predictable and secure. Enhanced tools and new capabilities in this release enable administrators to tailor the application environment to efficiently monitor and manage compute resources and security.

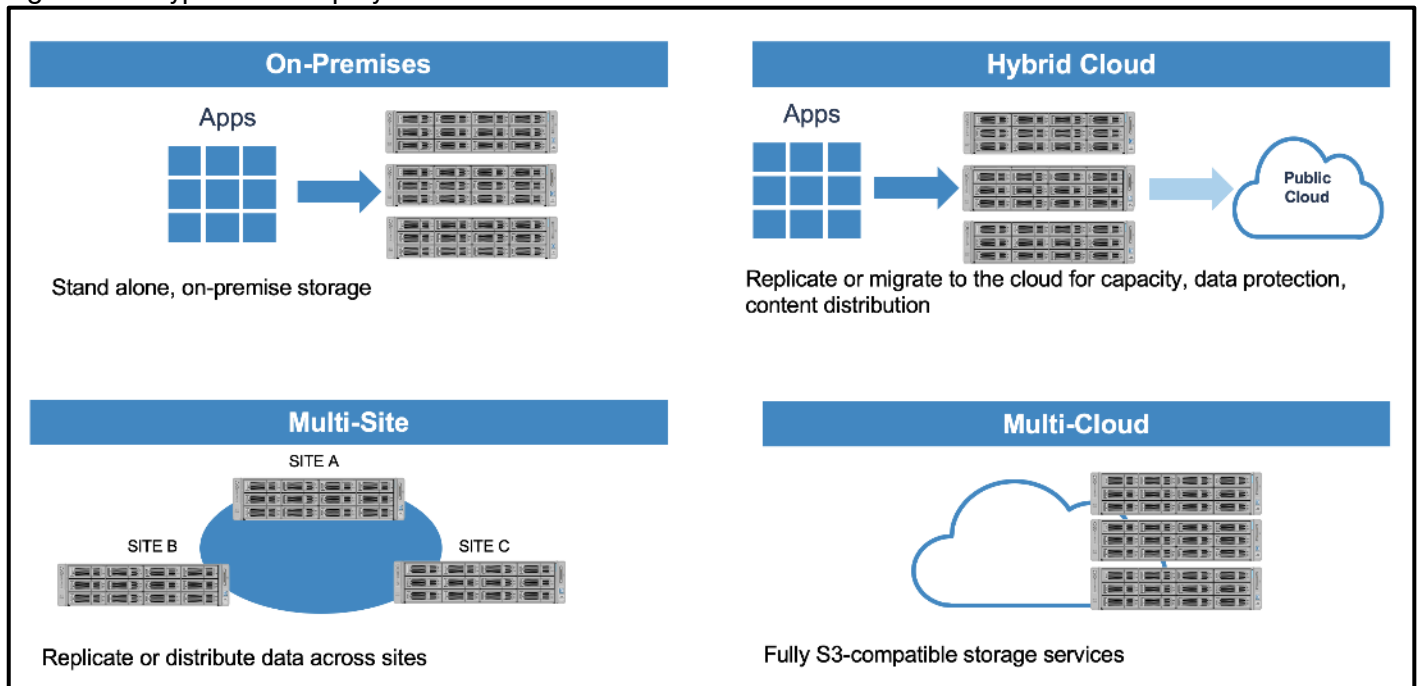
Clouidian HyperStore

Clouidian HyperStore enables data centers to provide highly cost-effective on-premise unstructured data storage repositories. Clouidian HyperStore is built on standard hardware that spans across the enterprise as well as into public cloud environments. Clouidian HyperStore is available as a stand-alone software. It easily scales to limitless capacities and offers multi-data center storage. HyperStore also has fully automated data tiering to all major public clouds, including AWS, Azure and Google Cloud Platform. It fully supports S3 applications and has flexible security options.

Clouidian HyperStore is a scale-out object storage system designed to manage massive amounts of data. It is an SDS solution that runs on the Cisco UCS platform allowing cost savings for datacenter storage while providing extreme availability and reliability.

HyperStore deployment models include on-premises storage, distributed storage, storage-as-a-service, or even other combinations (Figure 8).

Figure 8 HyperStore Deployment Models



Clouidian Object Storage

Clouidian deliver an object storage solution that provides petabyte-scalability while keeping it simple to manage. Deploy as on-premises storage or configure a hybrid cloud and automatically tier data to the public cloud.

View system health, manage users and groups and automate tasks with Clouidian’s web-based UI and REST API. Manage workload with a self-service portal that lets users administer their own storage. Powerful QoS capabilities help you ensure SLAs.

Clouidian makes it easy to get started. Begin with the cluster size that fits the needs and expand on demand. In Clouidian’s modular, shared-nothing architecture, every node is identical, allowing the solution to grow from a few nodes to a few hundred without disruption. Performance scales linearly, too.

Cloudian HyperStore offers a 100 percent native S3 API, proven to deliver the highest interoperability in its class. Guaranteed compatible with S3-enabled applications, Cloudian gives you investment protection and peace of mind.

Get the benefits of both on-premises and cloud storage in a single management environment. Run S3-enabled applications within data center with Cloudian S3 scale-out storage. Use policies you define to automatically tier data to the public cloud. It's simple to manage and limitlessly scalable.

Get all the benefits of using the Cisco UCS platform while managing data through a single pane of glass.

Cloudian HyperStore Design

Cloudian HyperStore is an Amazon S3-compliant multi-tenant object storage system. The system utilizes a “non-SQL” (NoSQL) storage layer for maximum flexibility and scalability. The Cloudian HyperStore system enables any service provider or enterprise to deploy an S3-compliant multi-tenant storage cloud.

The Cloudian HyperStore system is designed specifically to meet the demands of high volume, multi-tenant data storage:

- Amazon S3 API compliance. The Cloudian HyperStore system 100% compatible with Amazon S3's HTTP REST API. Customer's existing HTTP S3 applications will work with the Cloudian HyperStore service, and existing S3 development tools and libraries can be used for building Cloudian HyperStore client applications.
- Secure multi-tenancy. The Cloudian HyperStore system provides the capability to have multiple users securely reside on a single, shared infrastructure. Data for each user is logically separated from other users' data and cannot be accessed by any other user unless access permission is explicitly granted.
- Group support. An enterprise or work group can share a single Cloudian HyperStore account. Each group member can have dedicated storage space, and the group can be managed by a designated group administrator.
- Quality of Service (QoS) controls. Cloudian HyperStore system administrators can set storage quotas and usage rate limits on a per-group and per-user basis. Group administrators can set quotas and rate controls for individual members of the group.
- Access control rights. Read and write access controls are supported at per-bucket and per-object granularity. Objects can also be exposed via public URLs for regular web access, subject to configurable expiration periods.
- Reporting and billing. The Cloudian HyperStore system supports usage reporting on a system-wide, group-wide, or individual user basis. Billing of groups or users can be based on storage quotas and usage rates (such as bytes in and bytes out).
- Horizontal scalability. Running on standard off-the-shelf hardware, a Cloudian HyperStore system can scale up to thousands of nodes across multiple datacenters, supporting millions of users and hundreds of petabytes of data. New nodes can be added without service interruption.
- High availability. The Cloudian HyperStore system has a fully distributed, peer-to-peer architecture, with no single point of failure. The system is resilient to network and node failures with no data loss due to the automatic replication and recovery processes inherent to the architecture. A Cloudian HyperStore geocluster can be deployed across multiple datacenters to provide redundancy and resilience in the event of a data center scale disaster.

Cloudean HyperStore Architecture

The Cloudean HyperStore is a fully distributed architecture that provides no single point of failure, data protection options (replication or erasure coding), data recovery upon a node failure, dynamic re-balancing on node addition, multi-data center and multi-region support. Figure 9 illustrates the high-level system view.

Figure 9 High-level System View

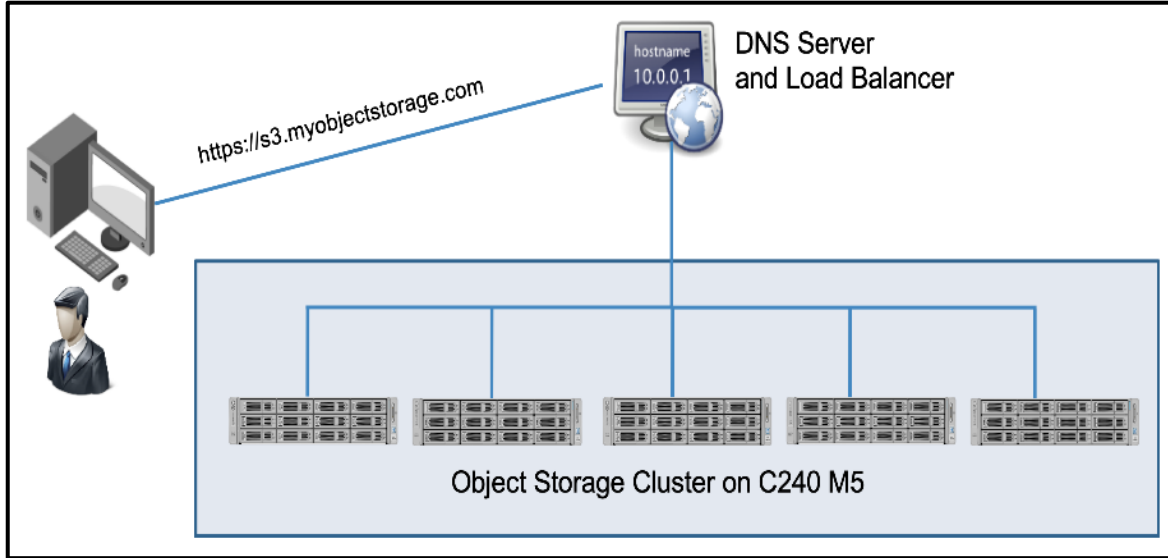
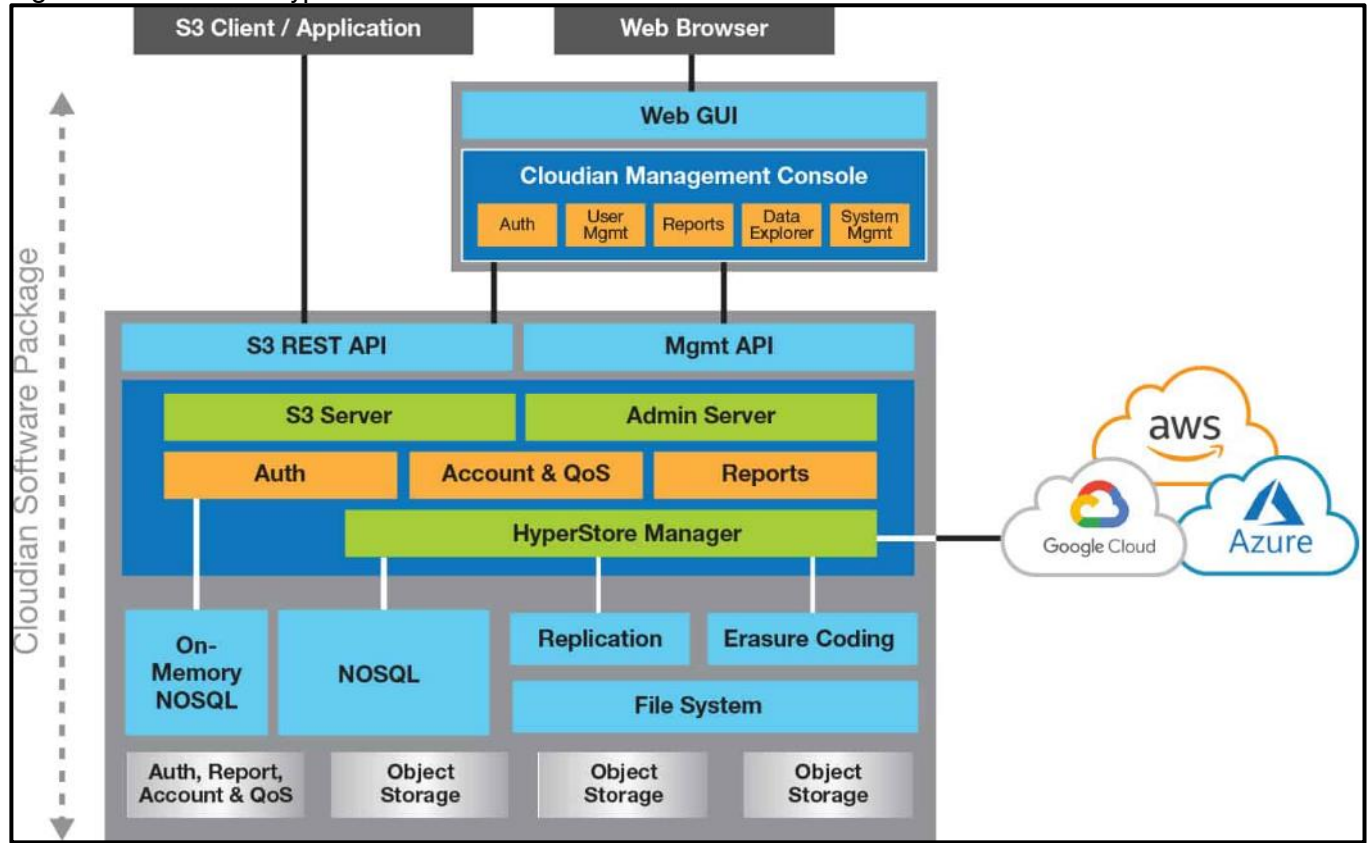


Figure 10 illustrates all of the service components that comprise a Cloudean HyperStore system.

Figure 10 Cloudian HyperStore Architecture



Cloudian Management Console

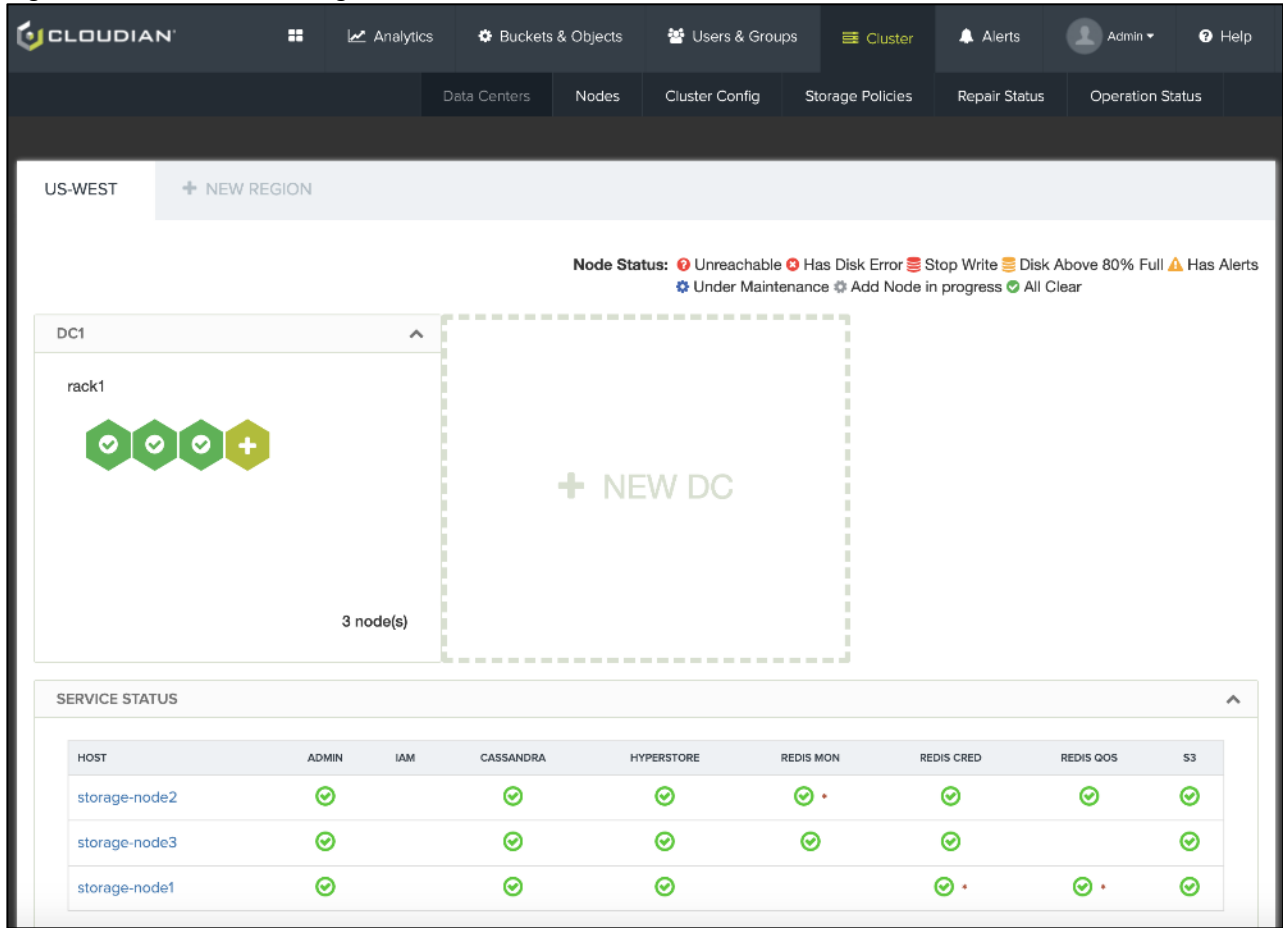
The Cloudian Management Console (CMC) is a web-based user interface for Cloudian HyperStore system administrators, group administrators, and end users. The functionality available through the CMC depends on the user type associated with a user's login ID (system administrative, group administrative, or regular user).

As a Cloudian HyperStore system administrator, you can use the CMC to perform the following tasks:

- Provisioning groups and users
- Managing quality of service (QoS) controls
- Creating and managing rating plans
- Generating usage data reports
- Generating bills
- Viewing and managing users' stored data objects
- Setting access control rights on users' buckets and stored objects

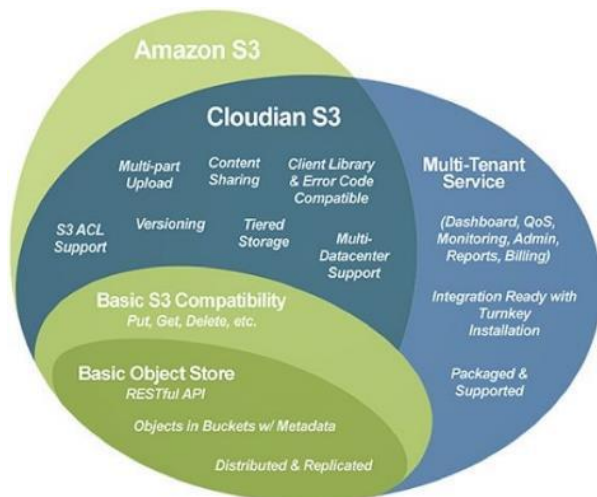
Group administrators can perform a limited range of administrative tasks pertaining to their own group. Regular users can perform S3 operations such as uploading and downloading S3 objects. The CMC acts as a client to the Administrative Service and the S3 Service.

Figure 11 Cloudian Management Console



S3 Compatible

With Amazon setting the cloud storage standard making it the largest object storage environment, and Amazon S3 API becoming the de facto standard for developers writing storage applications for cloud, it is imperative every Cloud, hybrid storage solution is S3 compliant. Cloudian HyperStore, in addition to being S3 compliant, also offers the flexibility to be on-premises object storage as well as hybrid tier to Amazon and Google clouds.



Integrated Billing, Management, and Monitoring

The HyperStore system maintains comprehensive service usage data for each group and each user in the system. This usage data, which is protected by replication, serves as the foundation for HyperStore service billing functionality. The system allows the creation of rating plans that categorize the types of service usage for single users or groups for a selected service period. The CMC has a function to display a single user’s bill report in a browser; HyperStore Admin API can be used to generate user or group billing data that can be ingested a third-party billing application. Cloudian HyperStore also allows for the special treatment of designated source IP addresses, so that the billing mechanism does not apply any data transfer charges for data coming from or going to these “whitelisted” domains.

ID	NAME	ACTIONS
Default-RP	Default Rating Plan	Edit
Whitelist-RP	Whitelist Rating Plan	Edit

Infinitely Scalable on Demand

Cisco and Cloudian HyperStore offers on-demand infinite scalability, allowing storage space to grow as needed. As demand grows, additional storage nodes can be added across multiple DCs.

Security

Cisco and Cloudian HyperStore takes safeguarding customer data very seriously. Two server-side encryption methods (SSE/SSE-c, KeySecure) are implemented to ensure that data is protected at all times.

Cloudian HyperStore simplifies the data encryption process by providing transparent key management at the server or node layer. This relieves administrators from the burden of having to manage encryption keys and eliminates the risk of data loss occurring due to lost keys. Furthermore, encryption can be managed very granularly—from a large-scale to an individual object.

Data Protection

With the ISA-L Powered Erasure Coding, Cloudian HyperStore optimizes storage for all data objects, providing efficient storage redundancy with low disk space consumption.

Effortless Data Movement

Cloudian HyperStore easily manages data, stores and retrieves data on-demand (with unique features like object streaming, dynamic auto-tiering), and seamlessly moves data between on-premises cloud and Amazon S3, irrespective of data size.

Solution Design

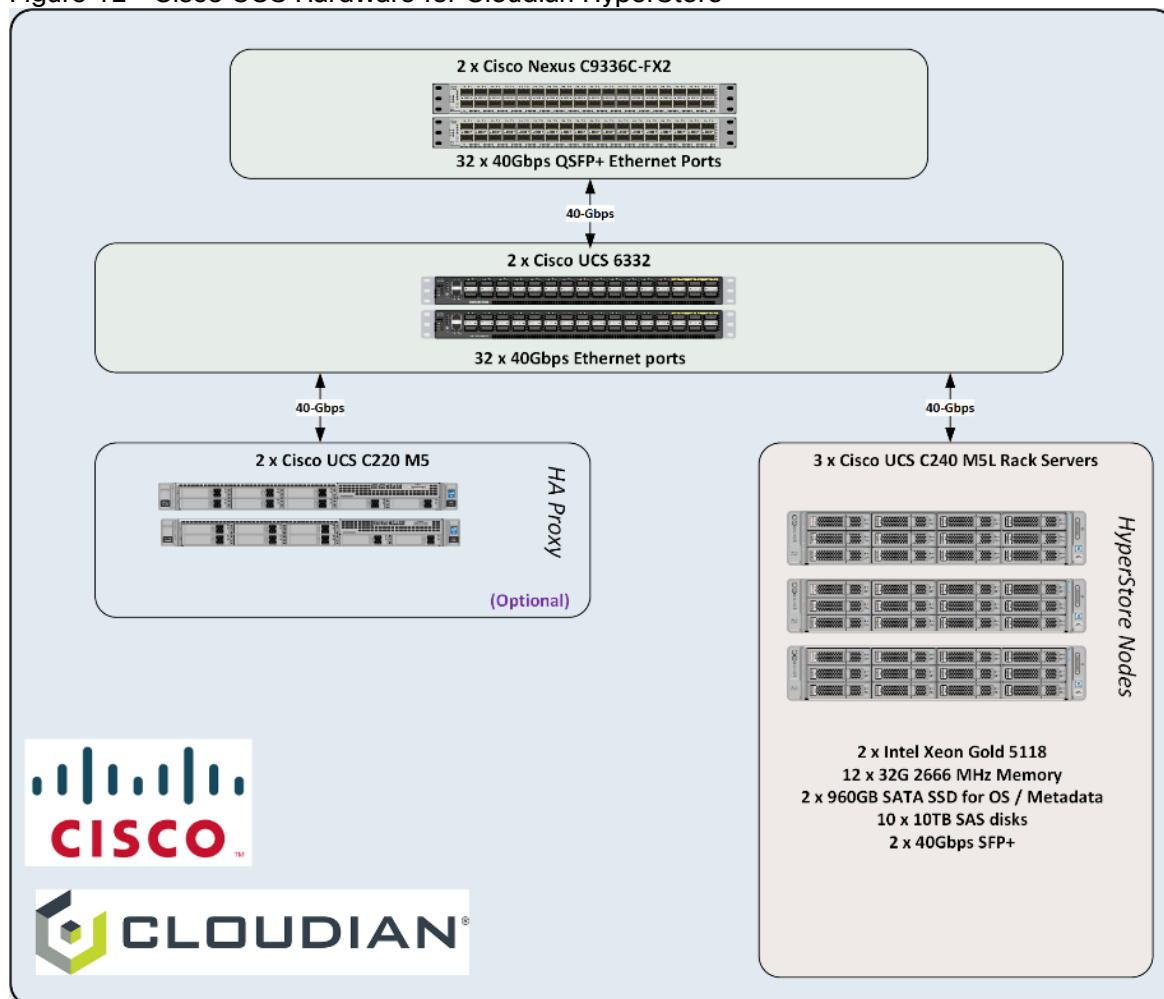
Deployment Architecture

The reference architecture use case provides a comprehensive, end-to-end example of designing and deploying Cloudian object storage on Cisco UCS C240 M5 as shown in Figure 12. This document describes the architecture and design of a Cloudian Scale-out object storage on three Cisco UCS C240 M5 Rack Servers and two Cisco UCS C220 M5S rack server as HA-proxy nodes. The whole solution is connected to a pair of Cisco UCS 6332 Fabric Interconnects and a pair of upstream network Cisco Nexus C9336C-FX2 switches.

The configuration is comprised of the following:

- 2 x Cisco Nexus 9000 C9336C-FX2 Switches
- 2 x Cisco UCS 6332 Fabric Interconnects
- 3 x Cisco UCS C240 M5L Rack Servers
- 2 x Cisco UCS C220 M5S Rack Servers (Optional for HA-Proxy)

Figure 12 Cisco UCS Hardware for Cloudian HyperStore



System Hardware and Software Specifications

Software Versions

Table 1 Software Versions

Layer	Component	Version or Release
Compute (Server/Storage Nodes)	BIOS	C240M5.4.0.4d.0.0506190827
	UCS C240 M5L	CIMC Controller 4.0(4c)
Compute (HA-Proxy Nodes)	BIOS	C220M5.4.0.4c.0.0506190754
	C220 M5S	CIMC Controller 4.0(4c)
Network	UCS Manager	4.0(4b)
	6332 Fabric Interconnect	Kernel 5.0(3)N2(4.04a) System 5.0(3)N2(4.04a)
Network	BIOS	05.33
Nexus 9000 C9336C-FX2	NXOS	9.2(3)
Software	Red Hat Enterprise Linux Server	7.6 (x86_64)
	Cloudian HyperStore	7.1.4

Hardware Requirements and Bill of Materials

Table 2 Bill of Materials

Component	Model	Quantity	Comments
Cloudian Storage Nodes	Cisco UCS C240 M5L Rack Servers	3	Per Server Node
			<ul style="list-style-type: none"> - 2 x Intel(R) Xeon(R) Gold 5118 (2.30 GHz/12 cores) - 384 GB RAM - Cisco 12G Modular Raid controller with 2GB cache - 2 x 960GB 3.5 inch Enterprise Value 6G SATA SSD (For OS and Metadata) - 10 x 10TB 12G SAS 7.2K RPM LFF HDD (512e) - Dual-port 40 Gbps VIC (Cisco UCS VIC 1385)
Cloudian HA-Proxy Node (Optional)	Cisco UCS C220 M5S Rack server	2	<ul style="list-style-type: none"> - 2 x Intel Xeon Silver 4110 (2.1GHz/8 Cores), 96GB RAM

Component	Model	Quantity	Comments
			<ul style="list-style-type: none"> - Cisco 12G SAS RAID Controller - 2 x 600GB SAS for OS - Dual-port 40 Gbps VIC
UCS Fabric Interconnects FI-6332	Cisco UCS 6332 Fabric Interconnects	2	
Switches Nexus 9000 C9336C-FX2	Cisco Nexus Switches	2	

Physical Topology and Configuration

Figure 13 illustrates the physical design of the solution and the configuration of each component.

The connectivity of the solution is based on 40 Gigabit. All components are connected via 40 QSFP cables. Between each Cisco UCS 6332 Fabric Interconnect and both Cisco Nexus 9336C-FX2 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.

Between both Cisco Nexus 9336C-FX2 switches are 4 x 40 Gbit cabling. Each Cisco UCS 6332 Fabric Interconnect is connected via 2 x 40 Gigabit to each Cisco UCS C9336C-FX2 switch. Cisco UCS C240 M5 and C220 M5 are connected via 1 x 40 Gbit to each Fabric Interconnect. The architecture is highly redundant, and system survived with little or no impact to applications under various failure test scenarios which will be covered in high availability testing.

Figure 13 Physical Design

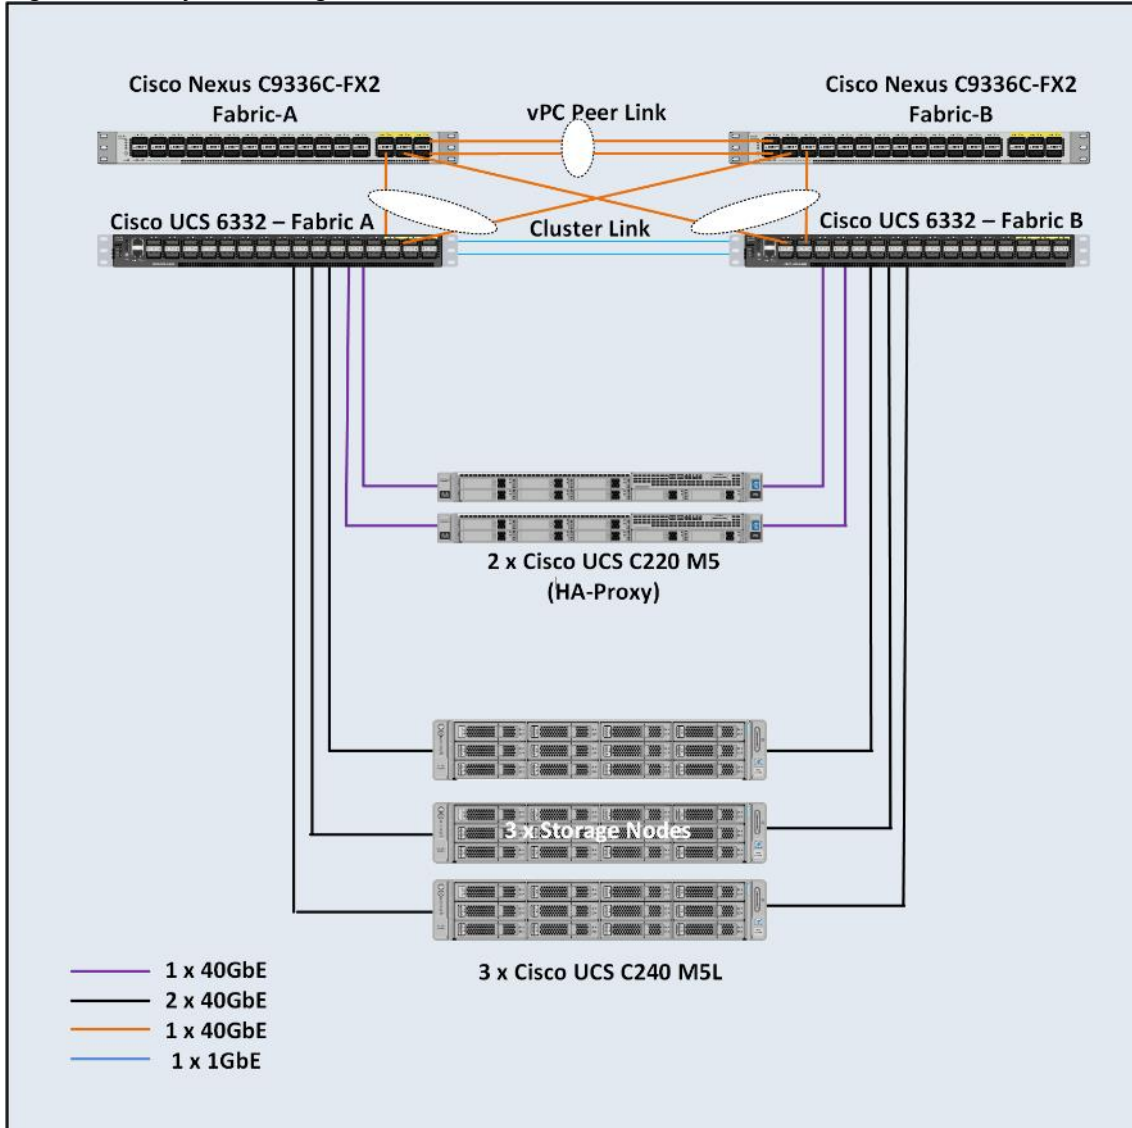
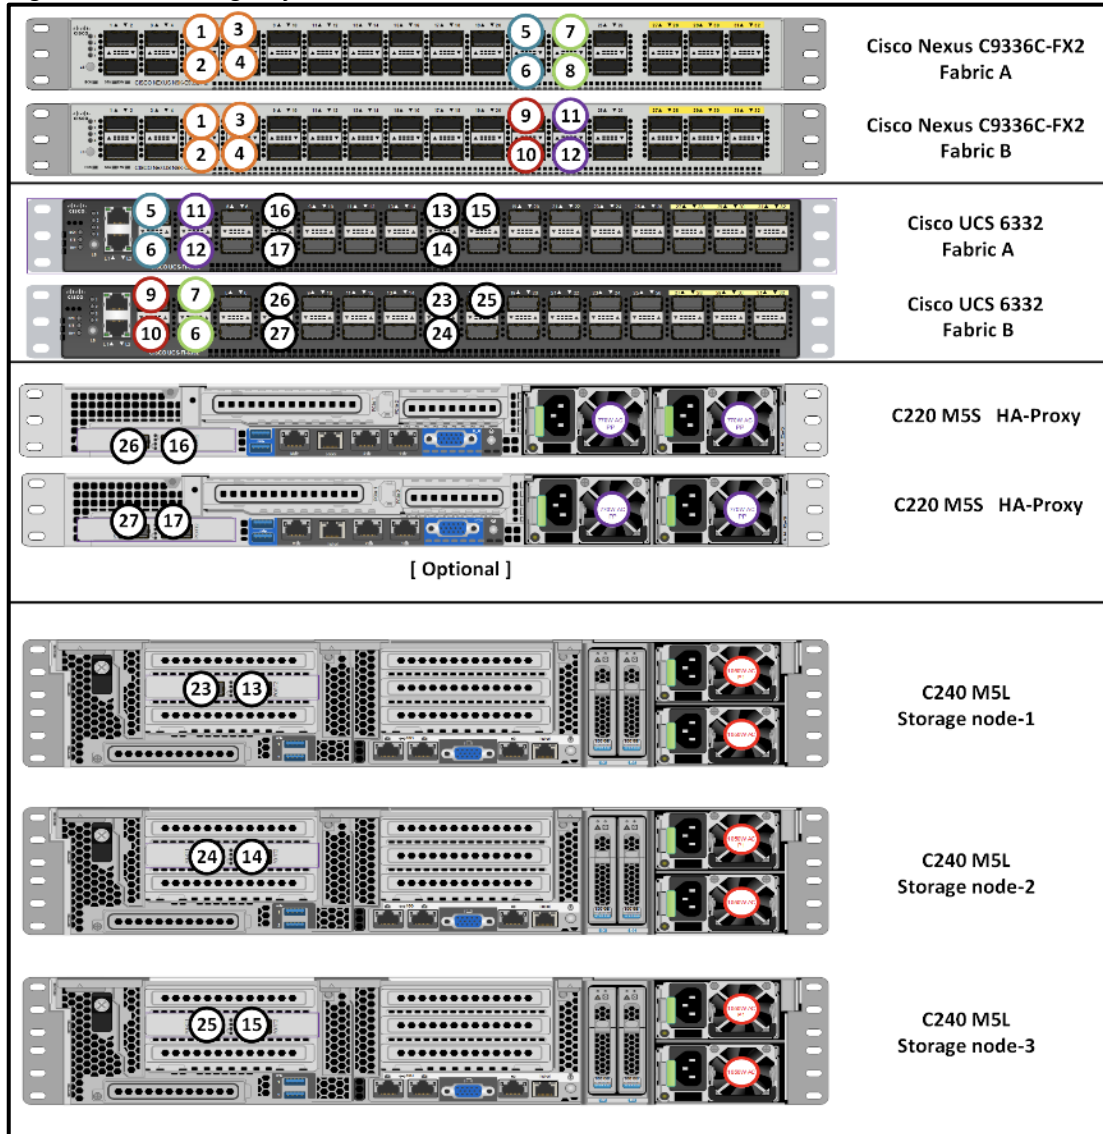


Figure 14 illustrates the actual cabling between servers and switches.

Figure 14 Cabling Layout

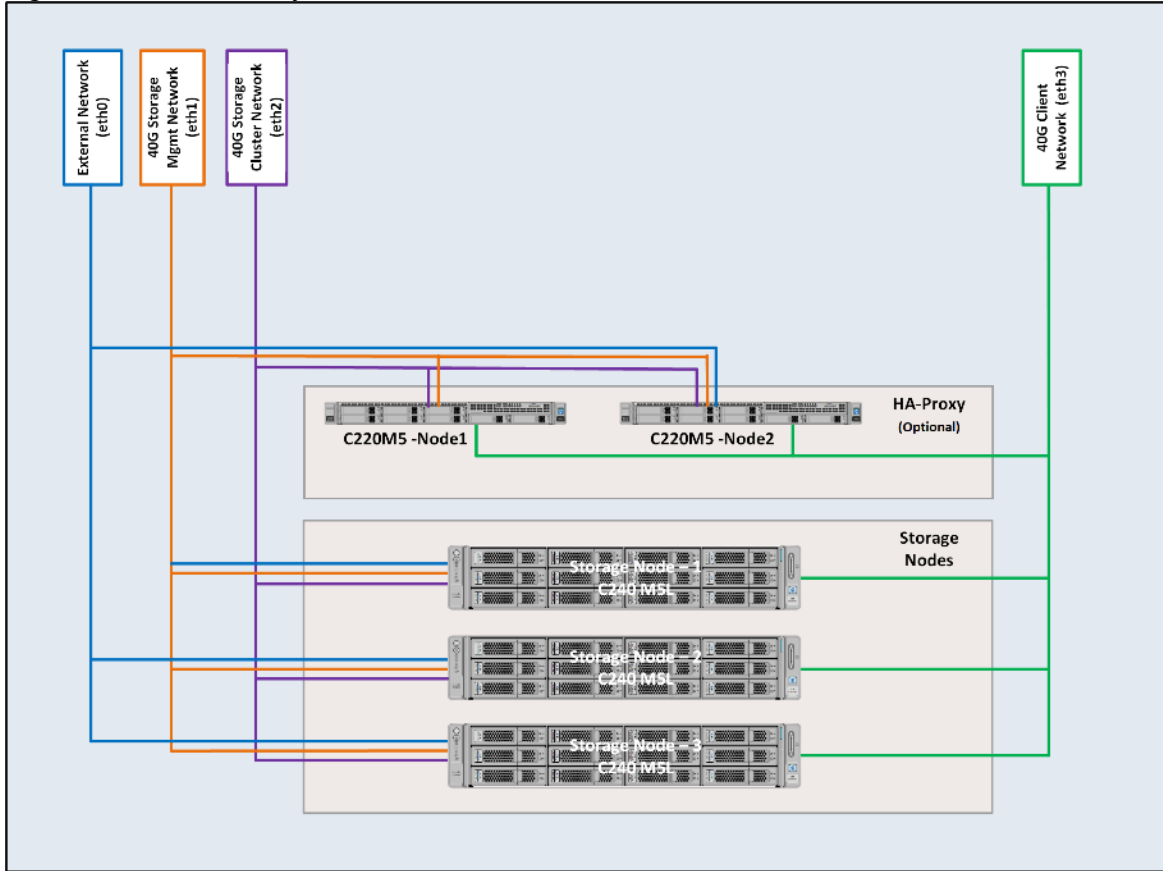


Network Topology

It is important to separate the network traffic with separate virtual NIC and VLANs for outward facing(eth0), host management(eth1), Cluster(eth2) and client(eth3) traffics. eth0, eth1 and eth3 are pinned to uplink interface 0 of VIC and eth2 is pinned to uplink interface 1 to enable better traffic distribution.

Figure 15 illustrates the Network Topology used in the setup.

Figure 15 Network Layout

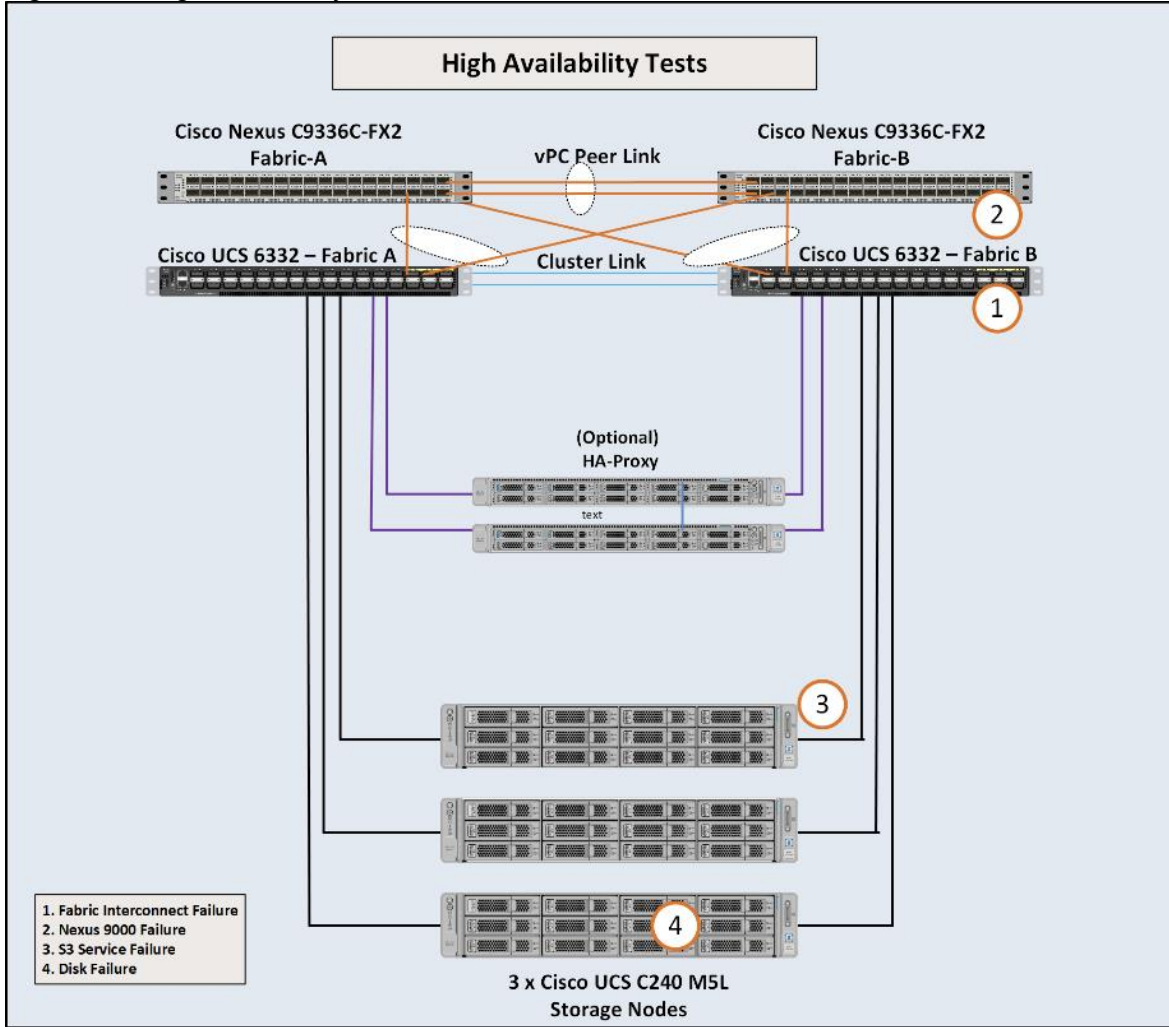


High Availability

As part of the hardware and software resiliency, random read and write load test with objects of 10MB in size will run during the failure injections. The following tests will be conducted on the test bed. The results of the tests will be included in this deployment guide.

1. Fabric Interconnect failures
2. Nexus 9000 failures
3. S3 Service failures
4. Disk failures

Figure 16 High Availability Tests



Design Criteria

There are several use cases and target industries where you can use the Cisco UCS and Cloudian SDS solution. The use cases and industries are several, but not limited to the following (Table 3).

Table 3 Use Cases

Primary	Backup and Archive
	Private & Hybrid Cloud
	Video/Content Distribution (VOD/Origin Server)
	Media Near-line Archive
	Medical Imaging
	Public Cloud - Email
	Public Cloud - Consumer Services
Secondary	Video Surveillance
	Enterprise File Sync and Share
	Hadoop Datalake
	Deep Learning

Table 4 Target Industries

Priority	Industry
Target	Telco, Mobile Operator & Cable Operator
	SaaS & Other Cloud Services
	Financial Services
	Media & Entertainment
	Police & Intelligence Agencies
	Hospitals and Medical Imaging Vendor
	Transportation
	Other Global 2000 (non XaaS, FIN,M&E,Transp,Hosp)

The following sections describe some considerations for the design of the Infrastructure and the Cloudian HyperStore.

Requirements

The requirements for the storage have to be understood for the design. These may include the total usable space, future expansion and organic growth for the capacity of the cluster, the performance of the cluster in terms of throughput and bandwidth, the average block size of IO, single site, multi-domain or multi-site requirements, and so on.

Design Considerations

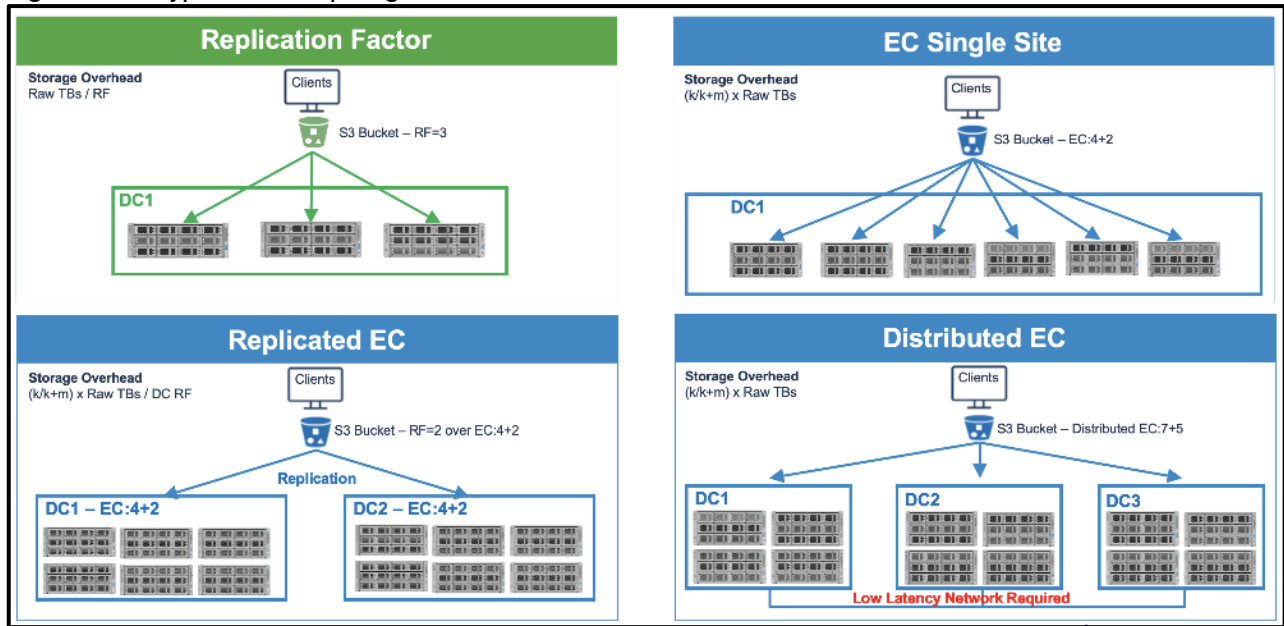
Number of Nodes of Cisco UCS C240 M5

When performance and storage capacity is not that important, a three-node configuration is recommended. This also reduces the TCO of the solution. However, as the performance and storage need increases, additional nodes can be added to the cluster.

Replication versus Erasure Coding

Central to Cloudbian’s data protection are its storage policies. These policies are ways of protecting data so that it’s durable and highly available to users. The Cloudbian HyperStore system lets you preconfigure one or more storage policies. Users, when creating a new storage bucket, can choose which preconfigured storage policy to use to protect data in that bucket. Users cannot create buckets until you have created at least one storage policy.

Figure 17 HyperStore Topologies



For each storage policy that you create, you can choose from the following two data protection methods:

Replication


With replication, a configurable number of copies of each data object are maintained in the system, and each copy is stored on a different node. For example, with 3X replication 3 copies of each object are stored, with each copy on a different node.

NUMBER OF DATACENTERS

1

DATA DISTRIBUTION SCHEME

Replicas Within Single Datacenter EC Within Single Datacenter



NUMBER OF REPLICAS

3

Erasure Coding


With erasure coding, each object is encoded into a configurable number (known as the “k” value) of data fragments plus a configurable number (the “m” value) of redundant parity fragments. Each fragment is stored on a different node, and the object can be decoded from any “k” number of fragments. For example, in a 4:2 erasure coding configuration (4 data fragments plus 2 parity fragments), each object is encoded into a total of 6 fragments which are stored on 6 different nodes, and the object can be decoded and read so long as any 4 of those 6 fragments are available.

NUMBER OF DATACENTERS

1

DATA DISTRIBUTION SCHEME

Replicas Within Single Datacenter EC Within Single Datacenter



ERASURE CODING K+M VALUE

- 4+2
- 6+2
- 8+2
- 9+3
- 12+4

Erasure coding requires less storage overhead (the amount of storage required for data redundancy) and results in somewhat longer request latency than replication. Erasure coding is best suited to large objects over a low latency network.

Supported Erasure Coding Configurations

Cloudian HyperStore supports EC, replicated EC, and distributed EC configurations.

- EC

This configuration requires a minimum 6 nodes across a single Data Centers (DC). This supports the minimum data and parity fragments of (4+2) where 2 is the parity. Table 5 lists the default EC configuration and the default number of nodes for a single DC.



Cloudian also supports 5 nodes EC as a custom policy – EC3+2.

Table 5 Default EC Configuration and Default Number of Nodes

Nodes in the DC	EC
6	4+2
8	6+2
10	8+2
12	9+3
16	12+4

- Replicated EC

This configuration requires a minimum of two Data Centers (DC). Each DC consists of 3 nodes each. This supports the minimum data and parity fragments of (2+1) where 1 is the parity. Table 6 lists the default replicated EC configuration and the default number of nodes per DC.

Table 6 Default Replicated EC Configuration and Default Number of Nodes

Nodes Total	DC1	DC2	EC
6	3	3	2+1
12	6	6	4+2
16	8	8	6+2
20	10	10	8+2
24	12	12	9+3

Each object is encoded into equal parts and parity fragments are replicated on each node. Each DC is a mirror image. For configurations greater than 2 DC, Distributed EC configuration is recommended. This configuration mirrors the encoded data and parity fragments to the other data centers in the configuration.

The choice among these three supported EC configurations is largely a matter of how many Cloudian HyperStore nodes in the datacenter. For a replicated EC configuration, a minimum of 3 nodes per DC are required.

- Distributed EC

Cloudian's Distributed EC solution implements the new ISA-L Erasure Codes that is vectored and fast. ISA-L is the Intel library containing functions to improve erasure coding.

The Cloudian Distributed Data Center with EC configuration requires a minimum of 3 data centers with 4 nodes each.

Data stored: DC1: 4, Dc2: 4, DC3:4, Metadata stored: Data stored: DC1: 4, DC2: 4, DC3:3

Distributed EC configuration offers the same level of protection as the replicated EC configuration with 50% less storage. The Distributed EC configuration is recommended if number of DCs involved are 3 or more.

Flash Storage

Flash Storage with SAS SSD's are used to store metadata for faster performance. The standard capacity requirement for Flash are less than 1 percent of the total data capacity. Standard design also calls for having a ratio of 1 SSD for 10 HDD.

JBOD versus RAID0 Disks

While Cloudian HyperStore as a SDS solution works with JBODs or with RAID0 disks, it is recommended to use JBOD for the solution. The 12G SAS RAID controller in C240 M5 provides up to 4G of cache that can be used for writes.

Memory Sizing

Memory sizing is based on the number of objects stored on each rack server, which is related to the average file size and the data protection scheme. Standard designs call for 384GB for the C240 M5.

Network Considerations

Cloudian Network requirements are standard Ethernet only. Please refer to the Network layout in Figure 10. While Cloudian software can work on a single network interface, it is recommended to create different virtual interfaces in Cisco UCS and segregate them. A client-access network and private-cluster network are required for the operation. Cisco UCS C240M5 has two physical ports of 40G each and the VIC allows you to create out many Virtual interfaces on each physical port.

It is recommended to have a private-cluster network on one port and the client-access networks on another port. This provides 40Gb bandwidth for each of these networks. While the client-access network requirements are minimal, every storage node can take up to 40Gb of client bandwidth requirements. Also, by having the client and cluster VIC's pinned to each fabric of the fabric interconnects, there is a minimal overhead of network traffic passing through the upstream switches for inter-node communication, if any. This unique feature of fabric interconnects and VIC's makes the design highly flexible and scalable.

Uplinks

The uplinks from fabric interconnects to upstream switches like Nexus, carry the traffic in case of FI failures or reboots. A reboot for instance is needed during a firmware upgrade. While there is complete high availability built-in the infrastructure, the performance may drop, depending on the uplink connectors from each FI to the Nexus vPC pool. If you want 'no' or a 'minimal drop', increase the uplink connectors.

Multi-Site Deployments

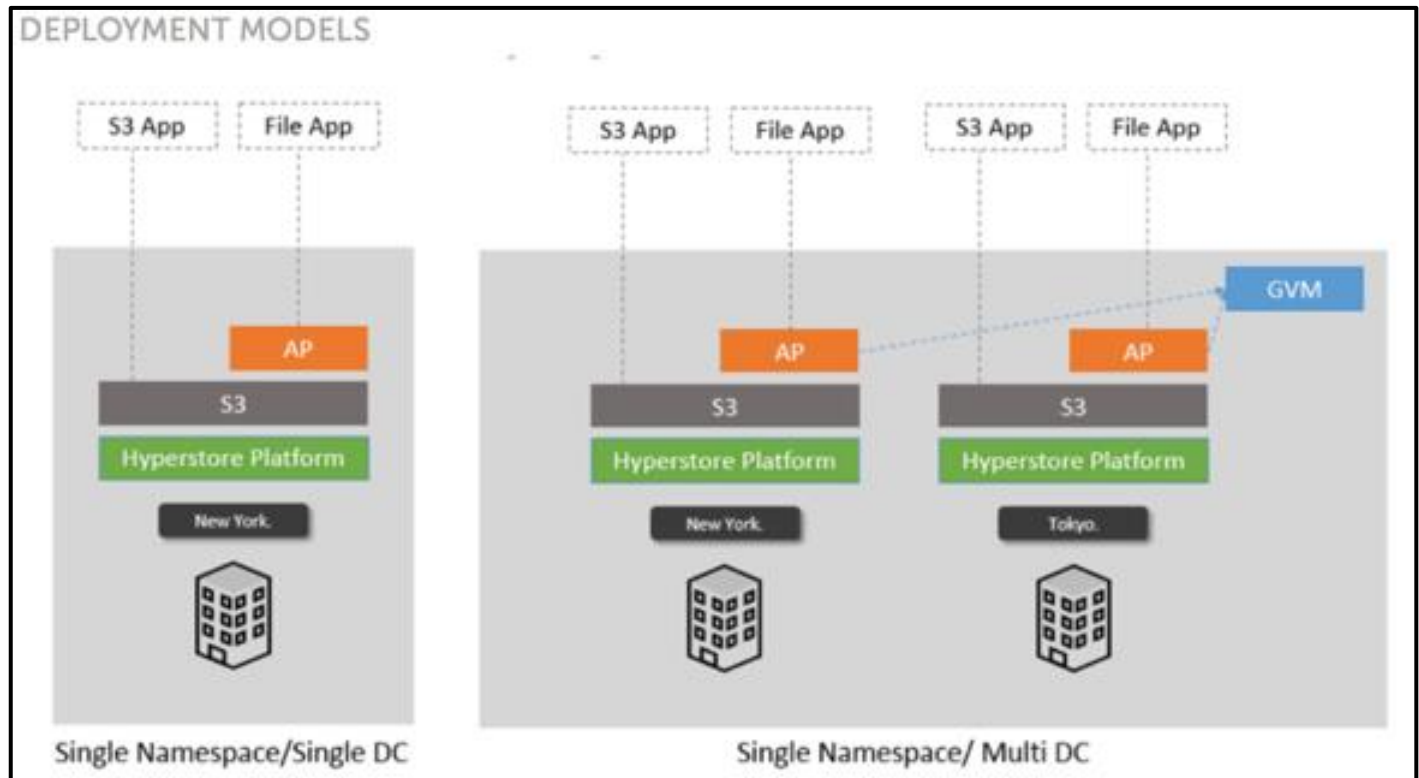
Similar to Amazon S3, the Cloudian HyperStore system supports the implementation of multiple "service regions." Setting up the Cloudian HyperStore system to use multiple service regions is optional.

The main benefits of deploying multiple service regions are:

- Each region has its own independent Cloudian HyperStore geo-cluster for S3 object storage. Consequently, deploying multiple regions is another means of scaling-out overall Cloudian HyperStore service offering (beyond using multiple nodes and multiple datacenters to scale out a single geo-cluster). In

a multi-region deployment, different S3 datasets are stored in each region. Each region has its own token space and there is no data replication across regions.

- With a multi-region deployment, service users can choose the service region in which their storage buckets will be created. Users may choose to store their S3 objects in the region that's geographically closest to them; or they may choose one region rather than another for reasons of regulatory compliance or corporate policy.



Designing a multi-site is beyond the scope of this document and for simplicity, only a single site deployment test bed is setup. Please contact Cisco and Clouidn if you have multi-site requirements.



Should a customer's workload and use case requirements not conform to the assumptions made while building these standard configurations, Cisco and Clouidn can work together to build custom hardware sizing to support the customer's workload.

Expansion of the Cluster

Cisco UCS hardware, along with Clouidn HyperStore, offers exceptional flexibility in order to scale-out as storage requirements change:

- Cisco UCS 6332 Fabric Interconnects have 32 ports each. Each server is connected to either of the FI's. Leaving the uplinks and any other clients directly connected to the Fabrics, 24-28 server nodes can be connected to FI pairs. If more servers are required, you should plan for a multi-domain system.

- Cisco UCS offers KVM management both in-band and out-of-band. In case out-of-band management is planned, you may have to reserve as many free IP's as needed for the servers. Planning while designing the cluster makes expansion very straightforward.
- Cisco UCS provides IP pool management, MAC pool management along with policies that can be defined once for the cluster. Any future expansion for adding nodes and so on, is just a matter of expanding the above pools.
- Cisco UCS is a template and policy based infrastructure management tool. All the identity of the servers is stored through Service Profiles that are cloned from templates. When a template is created, a new service profile for the additional server, can be created and applied on the newly added hardware. Cisco UCS makes Infrastructure readiness, extremely simple, for any newly added storage nodes. Rack the nodes, connect the cables, and then clone and apply the service profile.
- When the nodes are ready, you may have to follow the node addition procedure per the Cloudian documentation.

The simplified management of the infrastructure with Cisco UCS and well-tested node addition from Cloudian makes the expansion of the cluster very simple.

Performance Characterization

The solution was deployed with 3 x Cisco UCS C240 M5L nodes. While a few functional tests were completed, more in-depth tests will follow. The findings of these tests and the resulting best practices will be documented in the deployment guide, which will be published at a later date.

Below performance tests were run with a storage policy of replication factor of 3 (3 replicas within data center)

Read and write performance was recorded with Object sizes of 4KB, 16KB, 512KB, 1MB, 4MB and 10MB with different number of worker threads. (1, 10, 100, 200, 400 and 600).

Figure 18 Performance Snapshot of Read Bandwidth

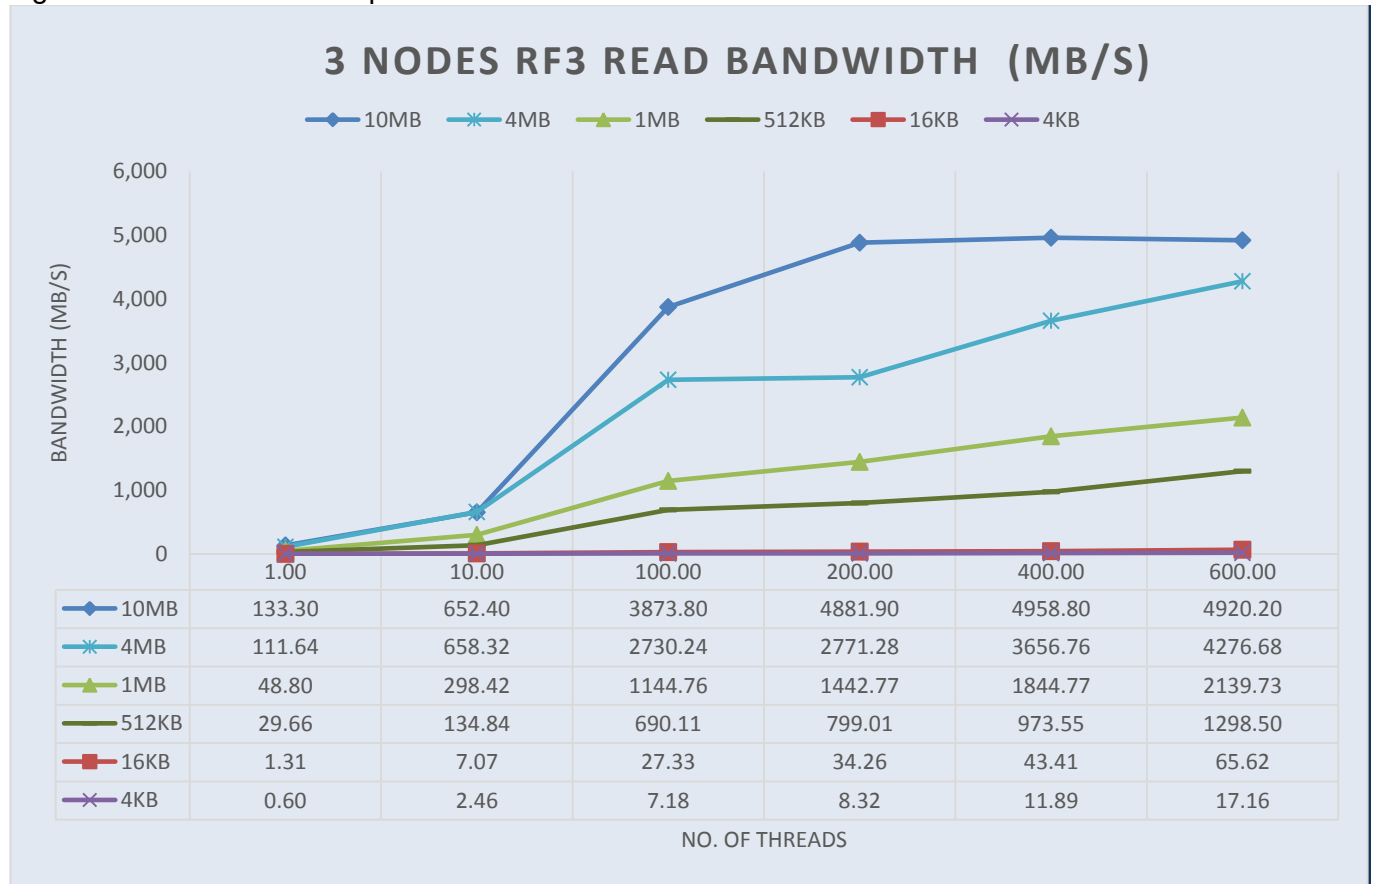


Figure 19 Performance Snapshot of Write Bandwidth

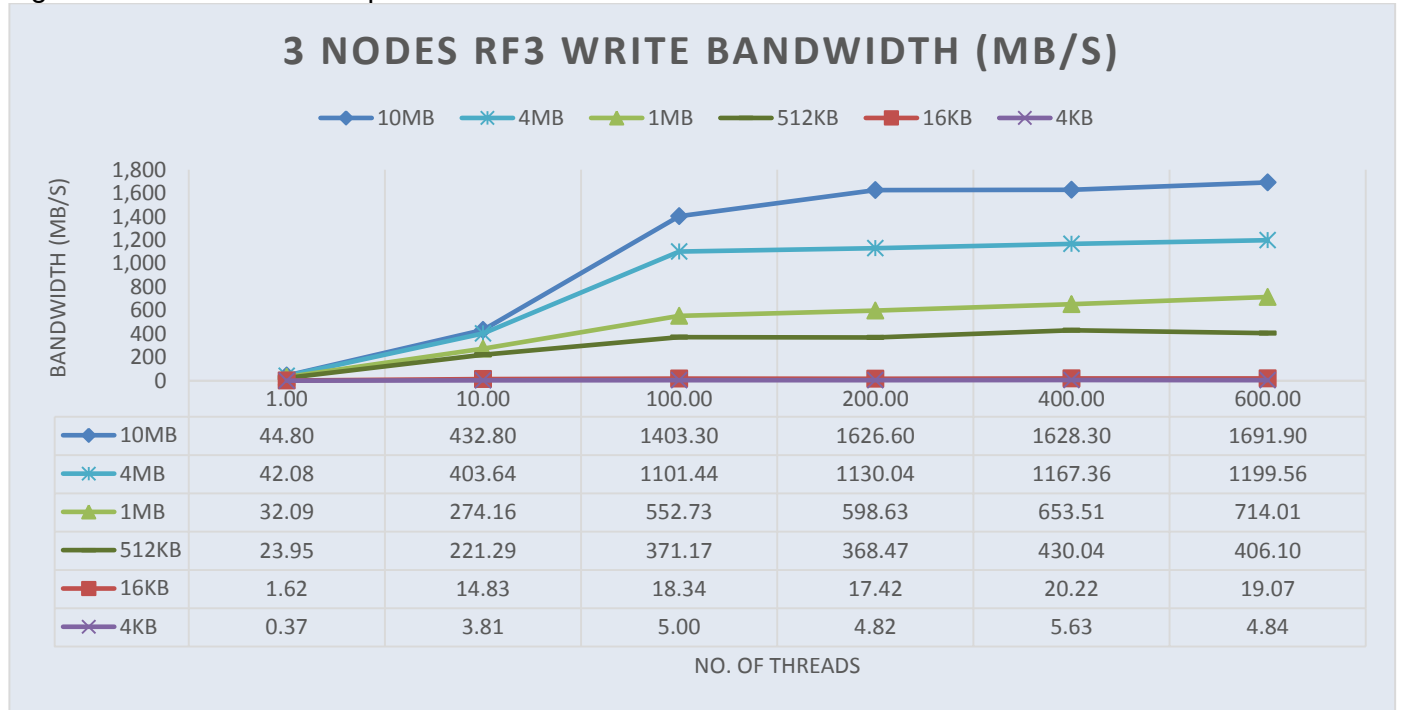


Figure 20 Performance Snapshot of Read IOPS




Figure 21 Performance Snapshot of Write IOPS




Figure 22 Performance Snapshot of Read Latency

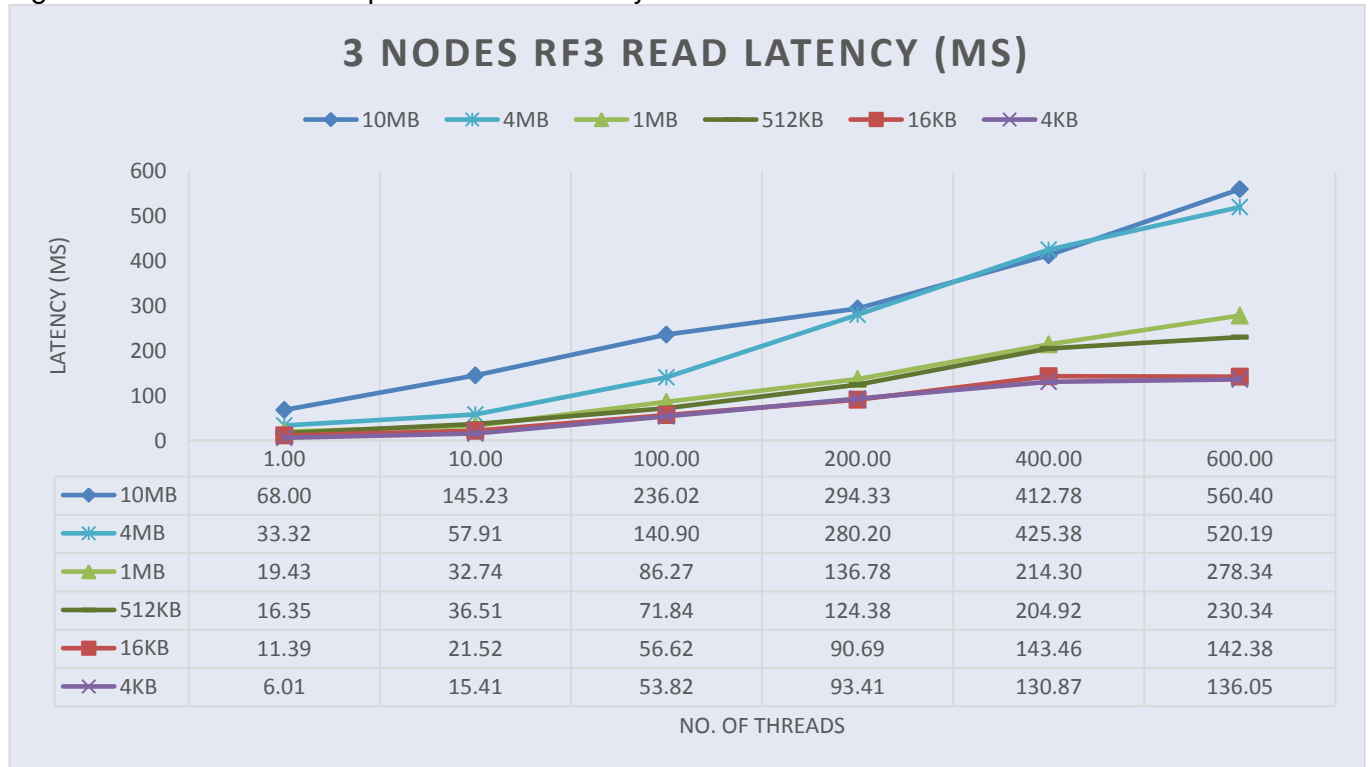
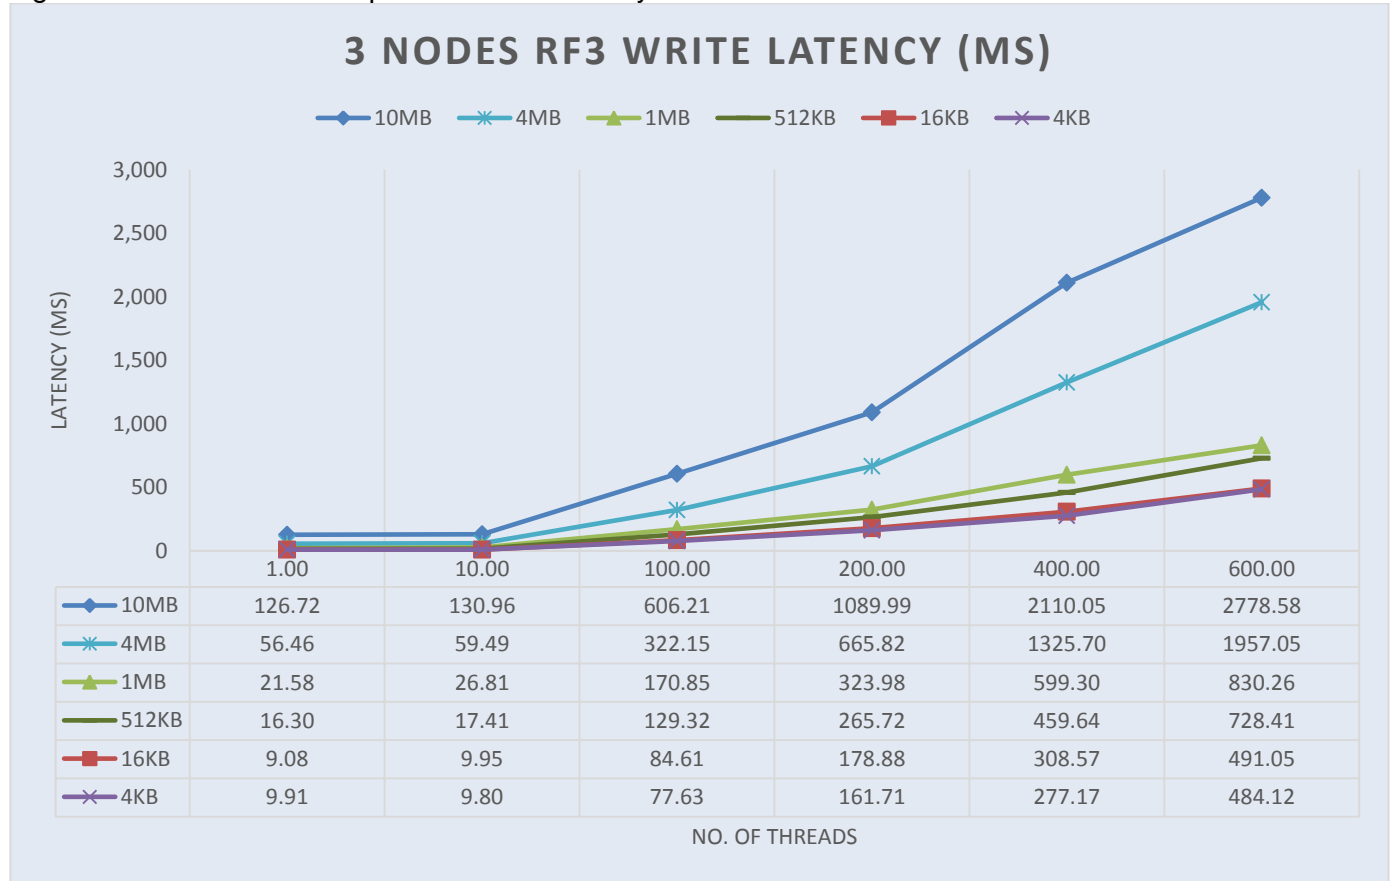


Figure 23 Performance Snapshot of Write Latency



Summary

Cisco UCS Infrastructure for Cloudian Software Defined Storage is an integrated solution to deploy Cloudian HyperStore and combines the value of Intel Xeon architecture, Cisco data center hardware and software, along with Red Hat Linux. This solution increases the speed of deployment and reduces the risk of scaling from proof-of-concept to full-enterprise production and is validated and supported by Cisco and Cloudian.

Cisco UCS hardware with Cisco UCS Manager Software brings an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Creating and cloning service profiles from its templates and maintaining the hardware from a single pane of glass not only provides rapid provisioning of hardware but also makes management and firmware upgrades simpler.

Cloudian HyperStore software makes it easy to build fully featured, Amazon S3-compliant cloud storage, on-premise. Cloudian HyperStore software ensures unlimited scale, multi-data center storage, fully automated data tiering, and support for all S3 applications—all behind firewall.

Cloudian HyperStore software deployed on Cisco UCS C240 M5L, combines robust availability with system management control, monitoring capabilities and reporting. A host of features, including hybrid cloud streaming, virtual nodes, configurable erasure coding, and data compression and encryption sets Cloudian apart with highly efficient storage and seamless data management. Built on a robust object storage platform for effortless data sharing, cloud service providers around the world use Cloudian HyperStore to deploy and manage both public and private clouds, while enterprises rely on it to maintain their private and hybrid clouds.

This Cisco Validated Design is a partnership of Cisco Systems, Inc. and Cloudian. Combining these technologies, expertise and experience in the field, we are able to provide an enterprise-ready hardware and software solution.

About the Authors

Paniraja Koppa, Cisco Systems, Inc.

Paniraja Koppa is a Technical Marketing Engineer for UCS Solutions. In his current role at Cisco, he works on best practices, optimization, automation and performance tuning of software defined storage solutions on Cisco UCS platforms. He has more than 13 years of experience with a primary focus on data center technologies such as Servers, Storage, Operating systems, Automation, Virtualization and Cloud. Prior to this, he has led QA efforts for 4 new virtual adapter card's firmware and software features for Cisco UCS. He also worked as customer advocate in the data center virtualization space.

Eddo Jansen, Cloudian, Inc

Eddo Jansen is Principal Architect at Cloudian. He has over 15 years of experience in IT Infrastructure, Storage, Virtualization and automation. His current role is building performant, scalable, highly available, and durable object store solutions with specialties in Performance testing, analyzing, troubleshooting and tuning.

Acknowledgements

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

- Chris O'Brien, Cisco Systems, Inc.
- Samuel Nagalingam, Cisco Systems, Inc.
- Muhammad Ashfaq, Cisco Systems, Inc.
- Oliver Walsdorf, Cisco Systems, Inc.
- Jawwad Memon, Cisco Systems, Inc.
- Sanjay Jagad, Cloudian