

Cisco UCS S3260 M5 Storage Server with Scality Software Defined Storage Design Guide

Last Updated: March 22, 2018



About the Cisco Validated Design (CVD) Program

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

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

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

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

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

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

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction	6
Audience	6
Purpose of this Document.....	6
Solution Summary.....	7
Technology Overview	8
Cisco Unified Computing System.....	8
Cisco UCS Manager	8
Cisco UCS 6300 Fabric Interconnects	9
Cisco UCS 9332 Nexus Switches	10
Cisco UCS S3260 M5 Storage Server.....	11
Cisco UCS C220 M5 Rack-Mount Server	13
Cisco UCS Virtual Interface Card 1387	13
Red Hat Enterprise Linux 7.4.....	14
Scality RING.....	15
Scale-Out Access Layer	16
RING Connectors	16
Local and Geo-Protection Layer	17
Scale-Out Object Storage Layer	17
Storage Nodes and IO Daemons.....	17
Scality S3 Connector	17
Scale-Out File System (SOFS)	19
Data Durability and Self-Healing	20
Replication Class of Service	20
Scality Erasure Coding.....	20
Self-Healing	21
Supervisor Web Management GUI	21
RING 7.4 New Features	22
S3 Connector Features.....	22
SOFS Connector Features	23
Management	23
Solution Design.....	24

Deployment Architecture	24
System Hardware and Software Specifications.....	25
Software Versions	25
Hardware Requirement and Bill of Materials.....	26
Physical Topology and Configuration	27
Network Topology	28
High Availability	29
Design Criteria	31
Requirements	31
Physical Infrastructure Considerations	32
Single Node versus Dual Node Cisco UCS S3260.....	32
Replication versus Erasure Coding.....	32
Flash Storage	32
JBOD versus RAID0 Disks.....	32
Memory Sizing.....	33
Network Considerations	33
Network Uplinks	33
Multi-Site Deployments	33
Connectors.....	33
Expansion of the Cluster	34
Deployment	35
Summary	37
About the Authors.....	38
Acknowledgements	38

Executive Summary

Cisco Validated Designs consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

The purpose of this document is to describe the design of Scality Ring on Red Hat Enterprise Linux and on latest generation of Cisco UCS S3260 M5 Servers. This validated design provides the framework of designing and deploying Scality SDS software on Cisco UCS S3260 storage servers. The Cisco Unified Computing System provides the storage, network, and storage access components for the Scality Ring, deployed as a single cohesive system.

Cisco Validated design describes how the Cisco Unified Computing System can be used in conjunction with Scality Ring 7.4. With the continuous evolution of SDS there has been increased demand to have Scality Ring validated on Cisco UCS servers. The Cisco UCS S3260 Storage Server, originally designed for the data center, together with Scality RING is optimized for object storage solutions, making it an excellent fit for unstructured data workloads such as backup, archive, and cloud data. The Cisco UCS S3260 delivers a complete infrastructure with exceptional scalability for computing and storage resources together with 40 Gigabit Ethernet networking.

Cisco and Scality are collaborating to offer customers a scalable object storage solution for unstructured data that is integrated with Scality RING Storage. 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.

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 is performed using **RESTful API's, which uses HTTP verbs such as GETs and PUTs**. Most of the newly generated data is unstructured today. With about 80 percent of data being unstructured, new approaches using x86 servers are proving to be more cost effective, providing storage that can be expanded as easily as your data grows. Scale-out Object storage is the newest approach for handling massive amounts of data.

The Scality RING is a Software-Defined Storage that is designed to create unbounded scale-out storage systems that converge the storage of Petabyte scale data from multiple applications and use-cases, including both object and file based applications.

Together with Cisco UCS, Scality Ring can deliver a fully enterprise-ready solution that can manage different workloads and still remain flexible. The Cisco UCS S3260 Storage Server is an excellent platform to use with the main types of Object and File workloads, such as capacity-optimized and performance-optimized workloads. It is best suited for sequential access, as opposed to random, to unstructured data, and to whatever the data size. It is essentially designed for Applications, not Direct end-users.

This document describes the architecture, design procedures of Scality storage on Cisco UCS S3260 M5 servers along with Cisco UCS C220 M5 rack-mount 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 UCS), Cisco Nexus and Cisco UCS Manager as well as a high-level understanding of Scality Ring 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 Scality Ring 7.4 on Cisco UCS platform. It discusses design choices and best practices using this shared infrastructure platform.

Solution Summary

This solution is focused on Scalify storage on Red Hat Linux 7.4 on Cisco Unified Computing System. The advantages of Cisco UCS and Scalify combine to deliver an object storage solution that is simple to install, scalable and performant. The configuration uses the following components for the deployment:

- Cisco Unified Computing System (Cisco UCS)
 - Cisco UCS 6332 Series Fabric Interconnects
 - Cisco UCS S3260 M5 storage servers
 - Cisco S3260 system IO controller with VIC 1380
 - Cisco C220M5 servers with VIC 1387
- Cisco Nexus C9332PQ Series Switches
- Scalify storage 7.4
- Red Hat Enterprise Linux 7.4

Technology Overview

Cisco Unified Computing System

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

The main components of Cisco Unified Computing System are:

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

The Cisco Unified Computing System is designed to deliver:

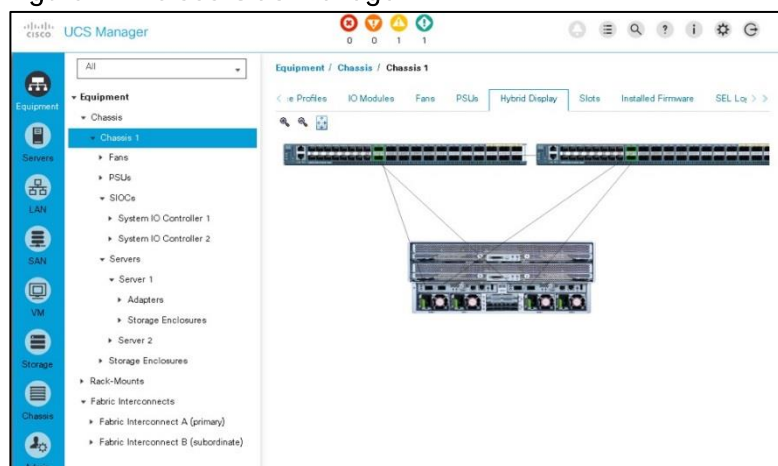
- A reduced Total Cost of Ownership (TCO) and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system which unifies the technology in the data center.
- Industry standards supported by a partner ecosystem of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

Cisco UCS Manager

Cisco UCS® Manager provides 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 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. Cisco Manager provides flexible role and policy-based management using service profiles and templates.

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

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

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 9332 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 9332 Switch



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

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

All 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 S3260 M5 Storage Server

The Cisco UCS® S3260 Storage Server 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 4 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® scalable processors, 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.

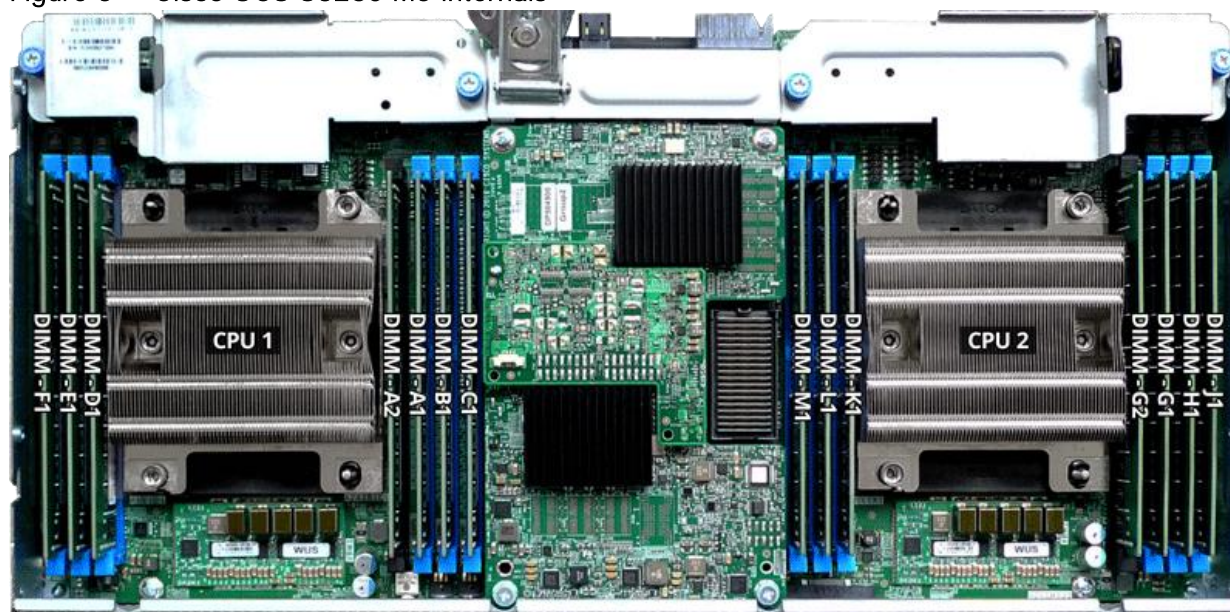
This high-density rack server comfortably fits in a standard 32-inch depth rack, such as the Cisco® R42610 Rack.

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 the following:

- Dual server nodes
- Up to 44 computing cores per server node
- Up to 60 drives mixing a large form factor (LFF) with up to 28 solid-state disk (SSD) drives plus 2 SSD SATA boot drives per server node
- Up to 1.5 TB of memory per server node (3 TB Total) with 128GB DIMMs
- Support for 12-Gbps serial-attached SCSI (SAS) drives
- A system I/O Controller either with HBA Passthrough or RAID controller, with DUAL LSI 3316 Chip
- 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
- Dual 7mm NVMe - Capacity points: 512G, 1TB and 2TB
- 1G Host Management Port

Figure 5 Cisco UCS S3260 M5 Internals



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 Unified Computing System™ 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.

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 your 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.

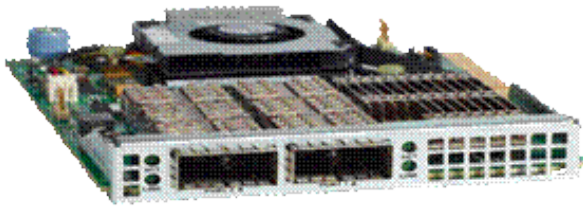
The Cisco UCS C220 M5 server included one dedicated internal modular LAN on motherboard (mLOM) slot for installation of a Cisco Virtual Interface Card (VIC) or third-party network interface card (NIC), without consuming a PCI slot, in addition to 2 x 10Gbase-T Intel x550 embedded (on the motherboard)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

The Cisco UCS C220 M5 server can be used standalone, or as part of 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 1387

The Cisco UCS Virtual Interface Card (VIC) 1387 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 3260 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 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.

Red Hat Enterprise Linux 7.4

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

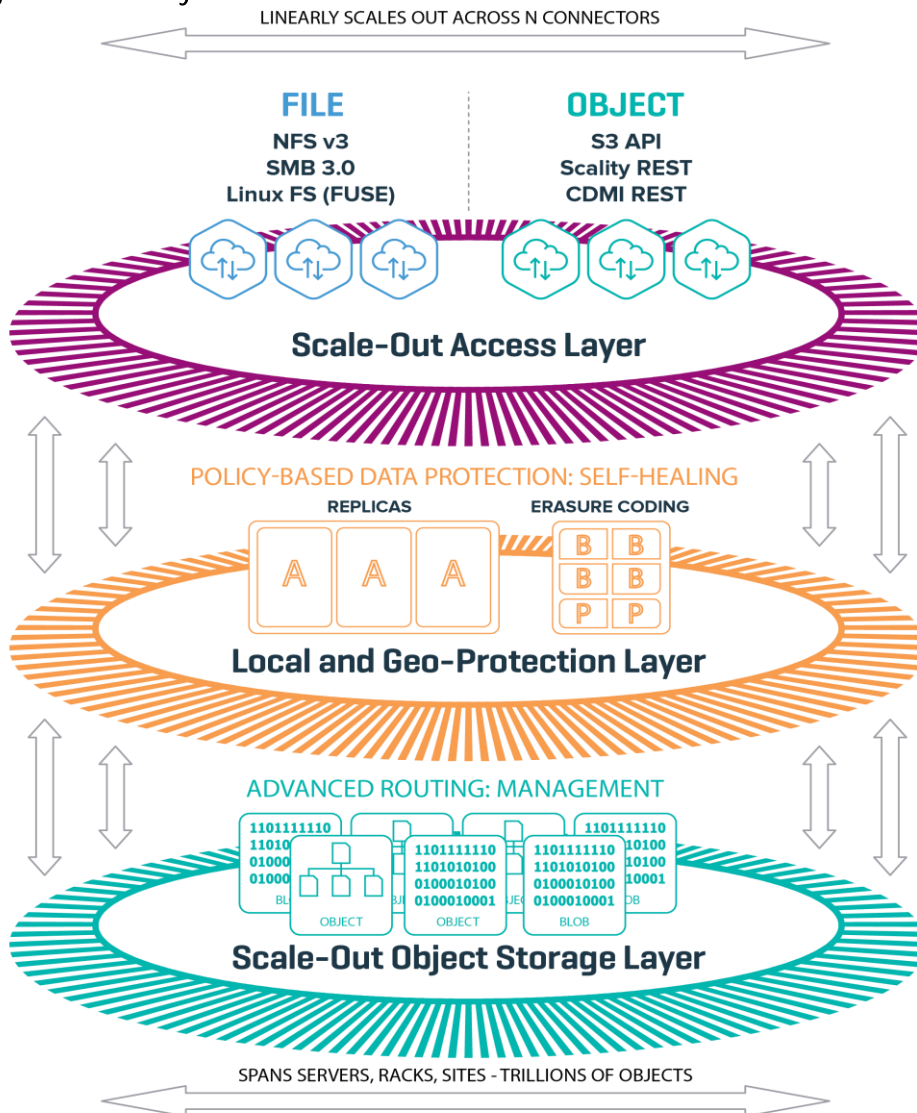
Designed to help organizations make a seamless transition to emerging datacenter models that include virtualization and cloud computing, Red Hat Enterprise Linux includes support for major hardware architectures, hypervisors, and cloud providers, making deployments across physical and different virtual 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.

Scality RING

The storage market has shifted dramatically in the last few years from one that is dominated by proprietary storage appliances. Scality RING is designed to support a broad variety of application workloads in a capacity-optimized fashion. The data center has evolved from providing mainly back-office transactional services, to providing a much wider range of applications including cloud computing, content serving, distributed computing and archiving.

Scality RING software is designed as a distributed, 100 percent parallel, scale-out architecture with a set of intelligent services for data access and presentation, data protection and systems management.

Figure 8 Scality RING Architecture



The RING is architected around 3 logical layers:

- Scale-out Access Layer

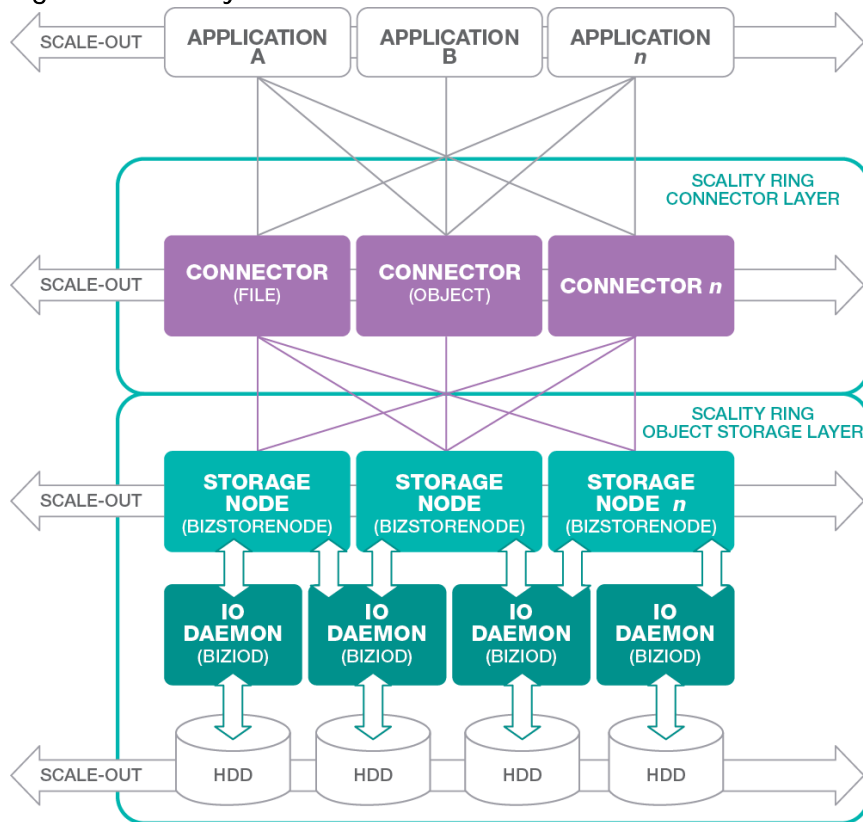
- Local and Geo Protection Layer
- Scale-out Object Storage Layer

At the heart of the storage layer is a distributed object key/value based store based on high performance peer to peer routing protocol.

Scale-Out Access Layer

The top layer has scalable access services (Connectors) that provide storage protocols for applications.

Figure 9 Scality Scale-Out Architecture



RING Connectors

The Connectors provide the top-level access points and protocol services for applications that use the RING for data storage. Applications may make use of multiple connectors in parallel to scale out the number of operations per second, or the aggregate throughput of the RING for high numbers of simultaneous user connections.

The RING Connectors provide a family of application interfaces including object-based Connectors, the S3 connector that is based on AWS S3, and Scality's native REST API, as well as file system Connectors (NFS v3, SMB 3.0, and FUSE) to fit a rich set of applications and a wide variety of data types.

Local and Geo-Protection Layer

This layer of the RING contains a set of data protection mechanisms to help ensure data durability and integrity, self-healing processes, and a set of systems management and monitoring services.

Data is persisted to the RING using two different protection schemes: Replication and Erasure Coding. These two models are described in more detail in a section below. All data protection is done extemporaneously and not after the fact. When an application receives acknowledgment that the data has been persisted, it means that it has been persisted already and fully protected.

Scale-Out Object Storage Layer

Storage Nodes and IO Daemons

Storage Nodes are virtual processes (Bizstorenode) that own and store a range of objects associated with its **portion of the RING's 'Keyspace'**. Each storage server is typically configured with six (6) storage node processes (Bizstorenode), and under these storage nodes are the storage IO daemons (Biziod), which are responsible for persistence of the data on disk, on a standard file system on a local disk. Each Biziod is a low-level process that manages the IO operations to a particular physical disk drive, maintaining the mapping between object ID's **and the actual object locations on disk**.

Each Biziod stores object payloads and metadata in a set of fixed size container files on each disk, with the storage daemon providing fast access for storage and retrieval operations into the container files. By containerizing objects, the system can still provide high-performance for small files and avoid the pitfalls of inode and management limits. The RING also leverages low-latency flash (SSD) devices for internal metadata for better performance.

The recommended deployment for systems that have both HDD and SSD media on the storage servers is to deploy a data RING on HDD, and the biziod metadata on SSD. Typically, the requirements for metadata are approximately 1 percent of the storage capacity of the actual data, so the sizing of SSD should follow that percentage for the best effect.

Scality S3 Connector

The Scality S3 Connector provides an advanced, modern S3 compatible application interface to the Scality RING. The AWS S3 service has become the leading cloud storage service and its API has emerged as the standard RESTful dialect for object storage, just like NFS was for the NAS generation. This is further amplified by the adoption of S3 among leading new and existing **ISV's who deliver solutions in areas such as** Backup and Archive (more traditional consumers of VTL and file-based storage interfaces), sync-n-share, file gateways and data mover solutions, media managers for video and image data and an increasing list of vertical industry solutions.

Rich AWS and Enterprise Security

Support for the full complement of AWS security services, such as multi-tenant accounts, Identity and Access Management (IAM) for users, groups and roles, AWS-style access keys and secret keys, the latest Signature v4 Authentication mechanism, and data encryption. Also featured is interoperability with such existing enterprise security services as LDAP and Microsoft® Active Directory® servers.

S3 API Compatibility

Notwithstanding rapid AWS advancements, a high-degree of S3 API coverage is assured, including core data APIs for Bucket and Object access and Multi-Part-Upload (MPU) for efficient ingest of large objects.

S3 Connector development is based on Continuous Integration (CI) and agile delivery of features when ready, which allows Scality to introduce new S3 methods shortly after their AWS publication. This functionality is provided by the S3 Server, which is supported by Scality as an open source project on GitHub.

Any-to-Any Scale-Out

Applications can access any Bucket or Object from any connector, thus allowing for parallel and multi-user access to data and scaling to billions of buckets and objects. Performance can be scaled-out simply by adding more connectors.

High-Performance Buckets

Support for low-latency response times and high throughput of reads and writes of Objects in Buckets. Also, performance is optimized for fast Bucket listing operations, including fast partial-path search for selected objects by path prefix.

Geo-Distributed Capabilities

S3 Connector provides integrated geo-replication capabilities for storage across multiple datacenters, supporting Active/Active stretched deployments for site disaster protection with continuous data availability and Site/Bucket level replication.

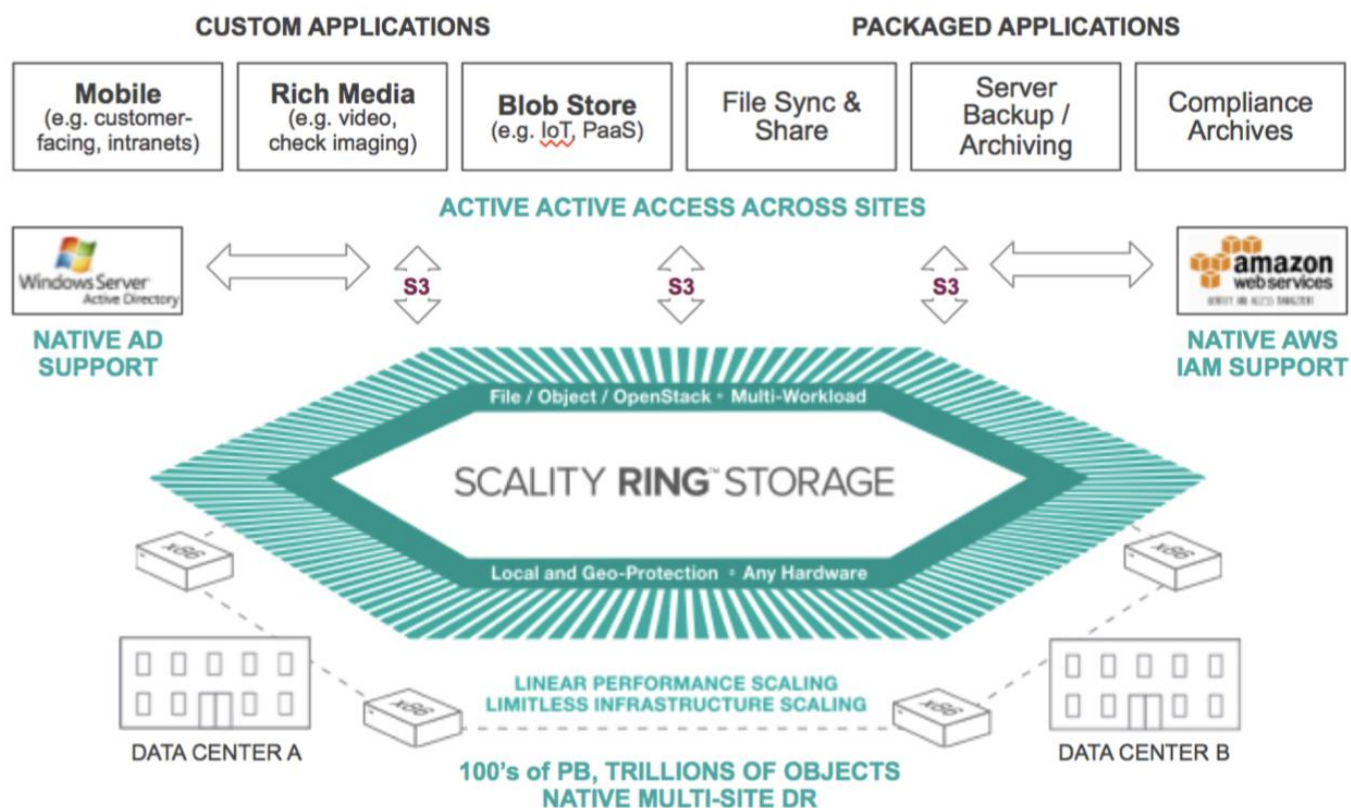
Object Versioning

Scality S3 Connector supports the AWS S3 Bucket Versioning API and follows the functional specifications of the AWS API. Versioning can be enabled and disabled on a per-Bucket basis. By enabling Versioning, the system will retain existing versions of an Object when it is modified. Previous versions of the object are therefore not overwritten but retained in a version history. Object reads will always access the most recent (current) version but can optionally specify a version ID to retrieve a specific older version of the object. This enables data restore capabilities if required in the event of a delete or inadvertent overwrite of the current version.

Ease of Deployment

Delivered as a set of easy-to-deploy Docker containers, installing the S3 Connector is simple, with zero-**configuration across the customer's choice of physical**, virtual or cloud environments.

Figure 10 Scality S3 Connector



Scale-Out File System (SOFS)

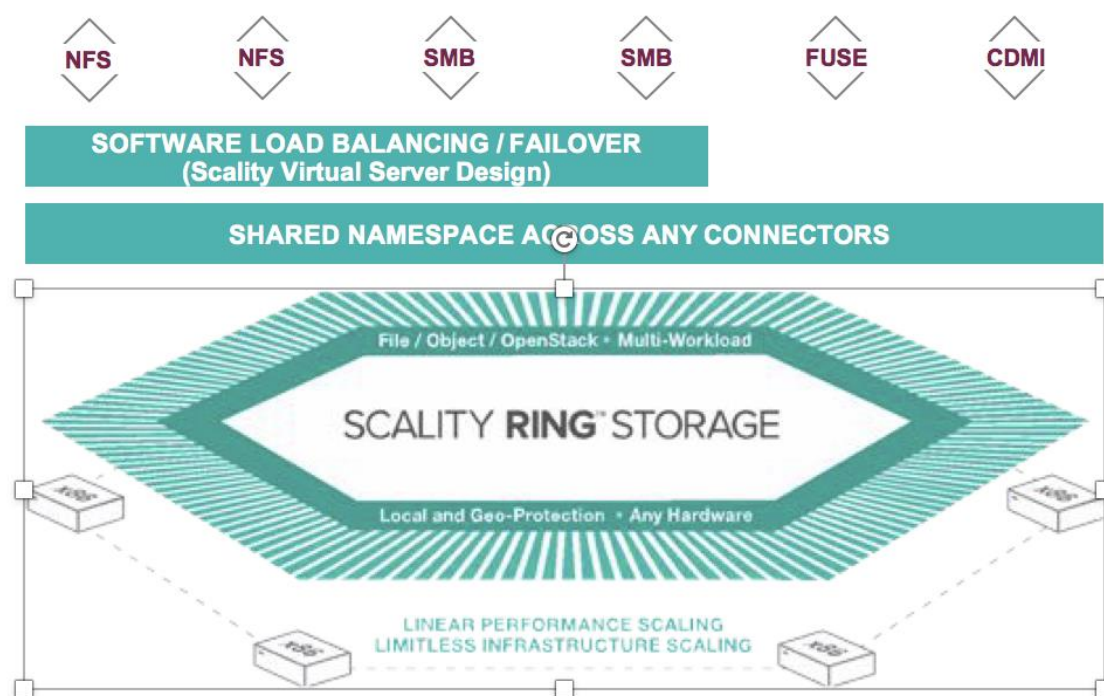
The RING supports native file system access to RING storage through the file Connectors and the integrated Scale-Out File System (SOFS). SOFS is a POSIX based virtual file system that provides file storage services without the need for external file gateways, as is common in other object storage solutions.

RING utilizes an **internal distributed database (MESA)** on top of the RING's storage services. MESA is a distributed database used to store file system directories and inode structures, providing virtual file system hierarchy, with guaranteed transactional consistency of file system data.

The RING provides the concept of “Volumes”, which may be used to easily configure file system services through the Supervisor. The RING can support up to 2^{32} volumes, with support for billions of files per volume.

- Connectors are stateless, can be IP load balanced, and do not lose or corrupt data if they fail
- File system state and INODES are stored in the highly available and durable RING using SSD for performance.

Figure 11 Scality SOFS with Single Namespace and Load Balancing



Data Durability and Self-Healing

The RING is designed to expect and manage a wide range of component failures including disks, servers, networks, and even across multiple data centers, while ensuring that data remains durable and available during these conditions. The RING provides data durability through a set of flexible data protection mechanisms optimized for distributed systems, including replication, erasure coding and geo-replication capabilities that allow applications to select the best data protection strategies for their data.

Replication Class of Service

To optimize data durability in a distributed system, the RING employs local replication, or the storage of multiple copies of an object within the RING. The RING will attempt to spread these replicas across multiple storage nodes, and across multiple disk drives, in order to separate them from common failures. While replication is optimal for many use cases where the objects are small, and access performance is critical, it does impose a high storage overhead penalty compared to the original data.

Scality Erasure Coding

Scality's Erasure Coding (EC) provides an alternative data protection mechanism to replication that is optimized for large objects and files. The basic idea with erasure coding is to break an object into multiple chunks (m), and apply a mathematical encoding to produce an additional set of parity chunks (k). The resulting set of chunks, $(m+k)$ are then distributed across the RING nodes, providing the ability to access the original object as long as any subset of m data or parity chunks are available. Stated another way, this provides a way to store an object with protection against k failures, with only k/m overhead in storage space.

Replication and EC may be combined, even on a single connector, by configuring a policy for the connector to store objects below a certain size threshold with replication, and files above with a specific EC schema. This allows the application to simply store objects without worrying about the optimal storage strategy per object, with the system managing that automatically.

Self-Healing

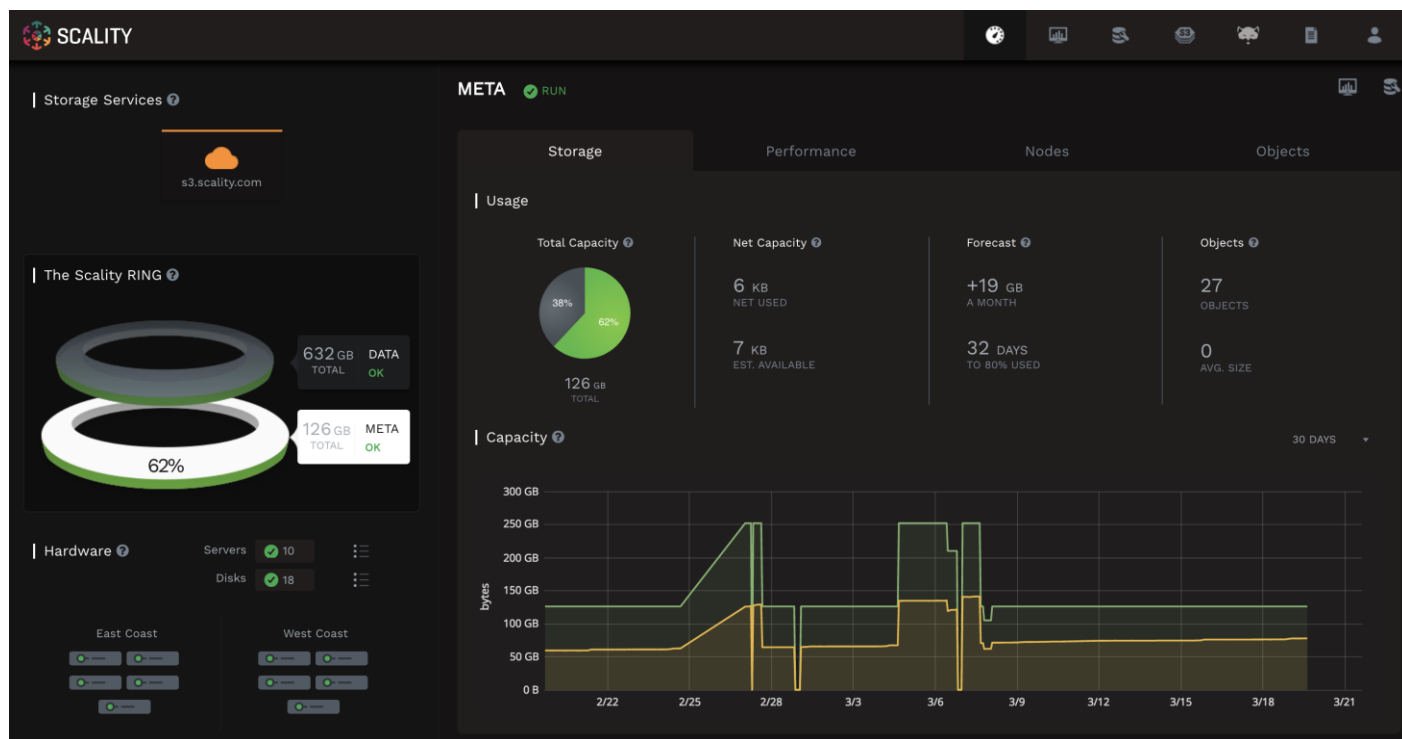
The RING provides self-healing operations to automatically resolve component failures, including the ability to rebuild missing data chunks due to disk drive or server failures, rebalance data when nodes leave and join the RING, and to proxy around component failures. In the event a disk drive or even a full server fails, background rebuild operations are spawned to restore the missing object data from its surviving replicas or EC chunks.

Supervisor Web Management GUI

The Supervisor is the RING's Web based management GUI. It provides visual, point-and-click style monitoring and management of the RING software, as well as the underlying physical platform layer. The Supervisor provides a main Dashboard page that provides graphical RING views, including the Servers, Zones and Storage Nodes comprising the RING, with browsing capabilities to drill down to details of each component, and pages for operations, management and provisioning of RING services. The Supervisor also provides performance statistics, resource consumption and health metrics through a rich set of graphs.

The Supervisor works in conjunction with the Scality management agent (*sagentd*), which is hosted on each Scality managed storage server, or connector server. The *sagentd* daemon provides a single point of communication for the Supervisor with the given host, for purposes of statistics and health metrics collection. This avoids the additional overhead of individual connections from the Supervisor to each Storage Node, and each disk drive daemon running on a specific host.

Figure 12 Supervisor Web GUI



RING 7.4 New Features

For a complete list of new features on RING 7.4 please refer to the Scality documentation located here: <http://www.scality.com>.

S3 Connector Features

- AWS S3 APIs support
- Vault - AWS Authentication (Signature v4 and v2)
- AWS S3 IAM Support
 - Groups
 - Policies
 - Roles
- Federated Access “Single Sign On” to S3 Connector
- Secure connections over HTTPS/SSL
- Bucket-Level Object Encryption with SafeNet KMS
- S3 Stretched deployments for 2 and 3-sites
- S3 CRR (Cross Regional Replication) for Asynchronous bucket replication
- S3 Bucket Service Utilization API (UTAPI) + Account level Utilization metrics

- IPv6 addresses on external connector interfaces
- Object versioning to track file revisions. The versioning functionality for S3 operations such as PUT, GET and DELETE requests is supported.
- Location Control for compliance and regulatory needs

SOFS Connector Features

- New SOFS Geo Models
 - 2-Site Stretched with Witness
 - 2-Site Asynchronous Replication
- GEOs Fail Back Improvements
- File Versioning and Versioning Policies
- Volume protection
- Access to File namespace through S3 API
- Enhanced logging (Volume-level space metering and quota)
- Access to File namespace through S3 API

Management

- Scality HALO Cloud Monitor
- Disk management tools
- New User Interfaces
 - Volume management and monitoring
 - S3 monitoring
- S3 Web browser
- Web S3 utilization per users/buckets
- Scality Cloud Monitor integration with the RING
- White branded Service Provider UI
- Disk management tools

Solution Design

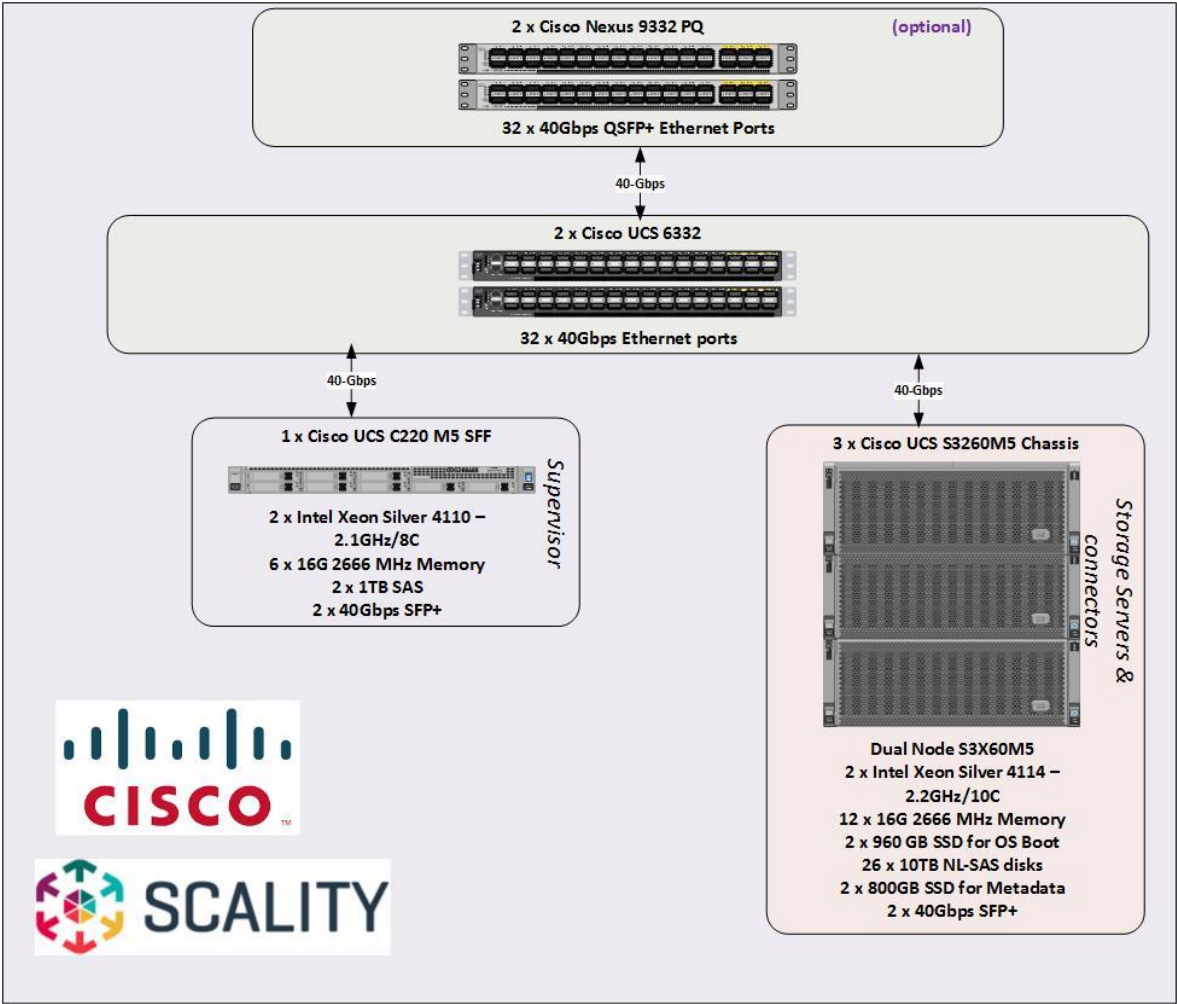
Deployment Architecture

The reference architecture use case provides a comprehensive, end-to-end example of designing and deploying Scality object storage on Cisco UCS S3260 as shown in Figure 13. This CVD describes the architecture and design of a Scality Scale-Out object Storage and file system solution on three Cisco UCS S3260 Storage Server Chassis', each with two Cisco UCS S3260 M5 nodes configured as Storage servers and one Cisco UCS C220 M5S Rack server as Supervisor node. The whole solution is connected to a pair of Cisco UCS 6332 Fabric Interconnects and to a pair of upstream network switches Cisco Nexus 9332PQ.

The detailed configuration is as follows:

- 2 x Cisco Nexus 9332PQ Switches
- 2 x Cisco UCS 6332 Fabric Interconnects
- 3 x Cisco UCS S3260 Storage Servers with 2 x Cisco UCS C3260 M5 server nodes each
- 1 x Cisco UCS C220 M5S Rack Servers

Figure 13 Cisco UCS Hardware for Scality RING



System Hardware and Software Specifications

Software Versions

Table 1 Software Versions

Layer	Component	Version or Release
Storage (Chassis) UCS S3260	Chassis Management Controller	3.1(3x)
	Shared Adapter	4.2(3x)
Compute (Server Nodes) UCS S3X60 M5	BIOS	S3260M5.3.1.3b
	CIMC Controller	3.1(3x)
Compute (Rack Server) C220 M5S	BIOS	C220M5.3.1.3b
	CIMC Controller	3.1(3x)
Cisco UCS 6332 Fabric	UCS Manager	3.2.3x

Layer	Component	Version or Release
Interconnect	Kernel	5.0(3)N2(3.23x)
	System	5.0(3)N2(3.23x)
Network Nexus 9332PQ	BIOS	07.59
	NXOS	7.0(3)I7(2)
Software	Red Hat Enterprise Linux Server	7.4 (x86_64)
	Scality RING	7.4

Hardware Requirement and Bill of Materials

Table 2 Bill of Materials

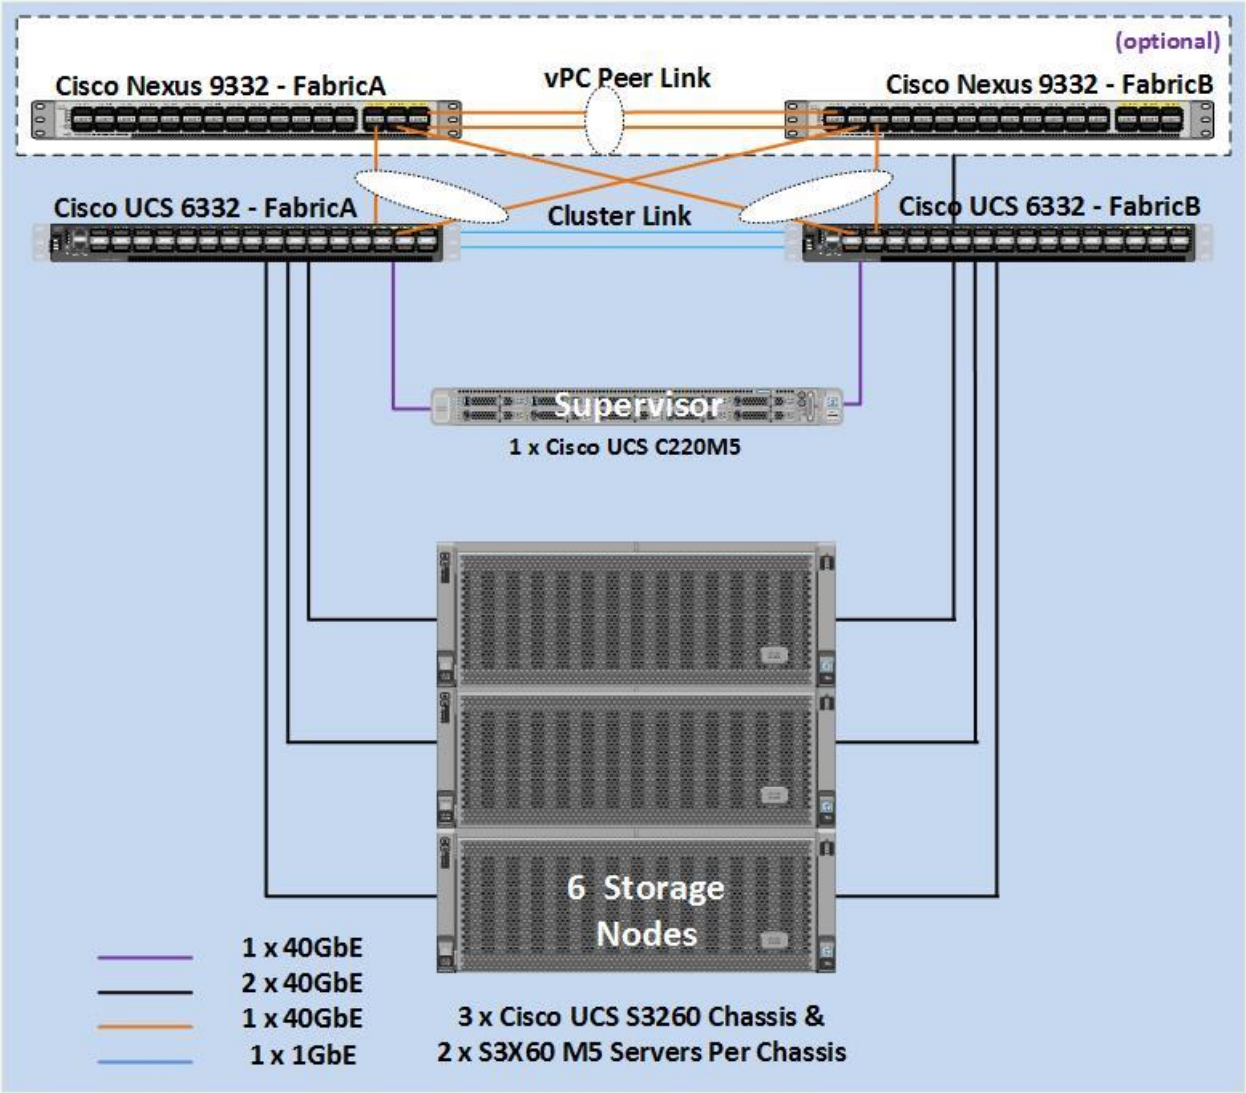
Component	Model	Quantity	Comments
Scality Storage Node (also the Connector Node)	Cisco UCS S3260 M5 Chassis	3	<ul style="list-style-type: none"> 2 x UCS S3X60 M5 Server Nodes per Chassis (Total = 6nodes) Per Server Node <ul style="list-style-type: none"> 2 x Intel Xeon Silver 4114 (2.2GHz/10cores), 192 GB RAM Cisco 12G SAS RAID Controller 2 x 1.6 TB SSD for OS 26 x 10TB HDDs for Data, 2 x 800G SSD for Metadata or 2 TB of NVMe's. Dual-port 40 Gbps VIC
Scality Supervisor Node	Cisco UCS C220 M5S Rack server	1	<ul style="list-style-type: none"> 2 x Intel Xeon Silver 4110 (2.1GHz/8 Cores), 96GB RAM Cisco 12G SAS RAID Controller 2 x 600GB SAS for OS Dual-port 40 Gbps VIC
UCS Fabric Interconnects	Cisco UCS 6332 Fabric Interconnects	2	
Switches	Cisco Nexus 9332PQ Switch-	2	

Component	Model	Quantity	Comments
	es		

Physical Topology and Configuration

The following sections describe the physical design of the solution and the configuration of each component.

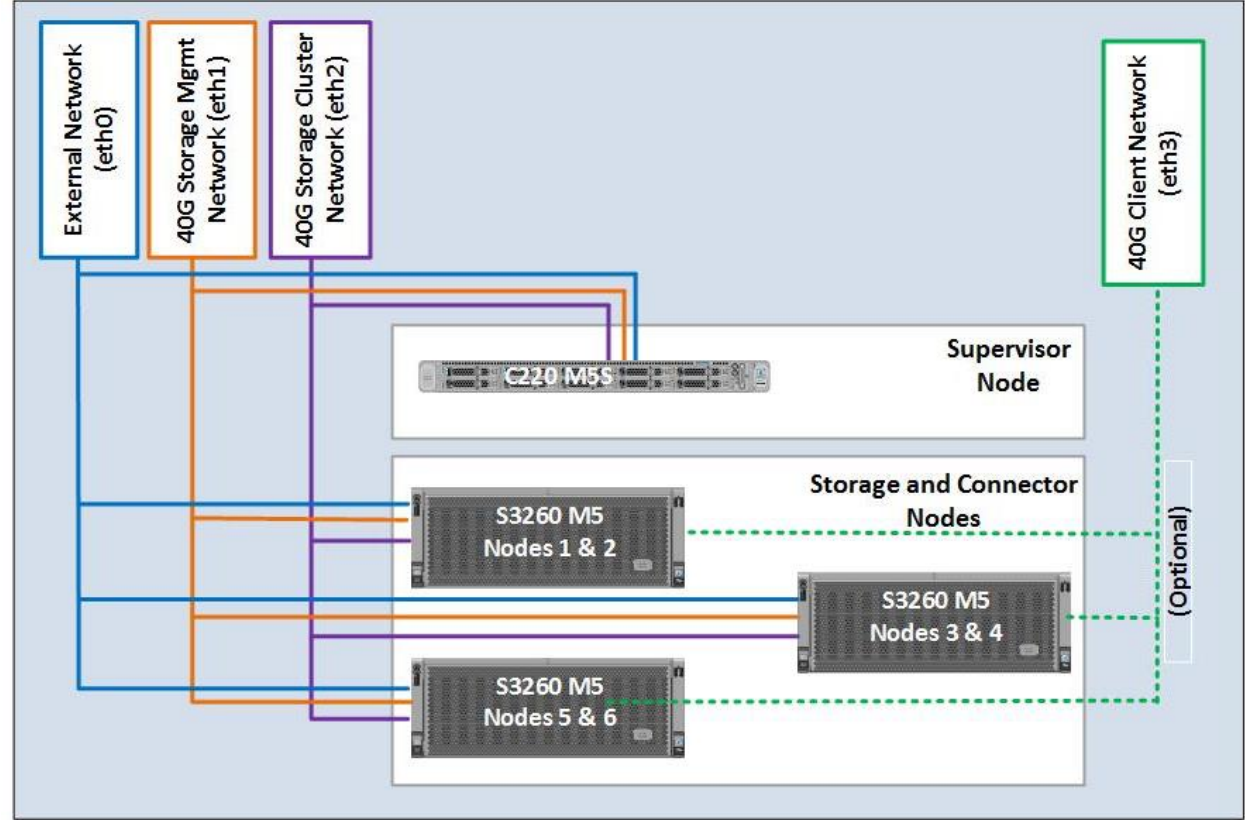
Figure 14 Physical Topology



Network Topology

Figure 15 illustrates the Network Topology used in the setup.

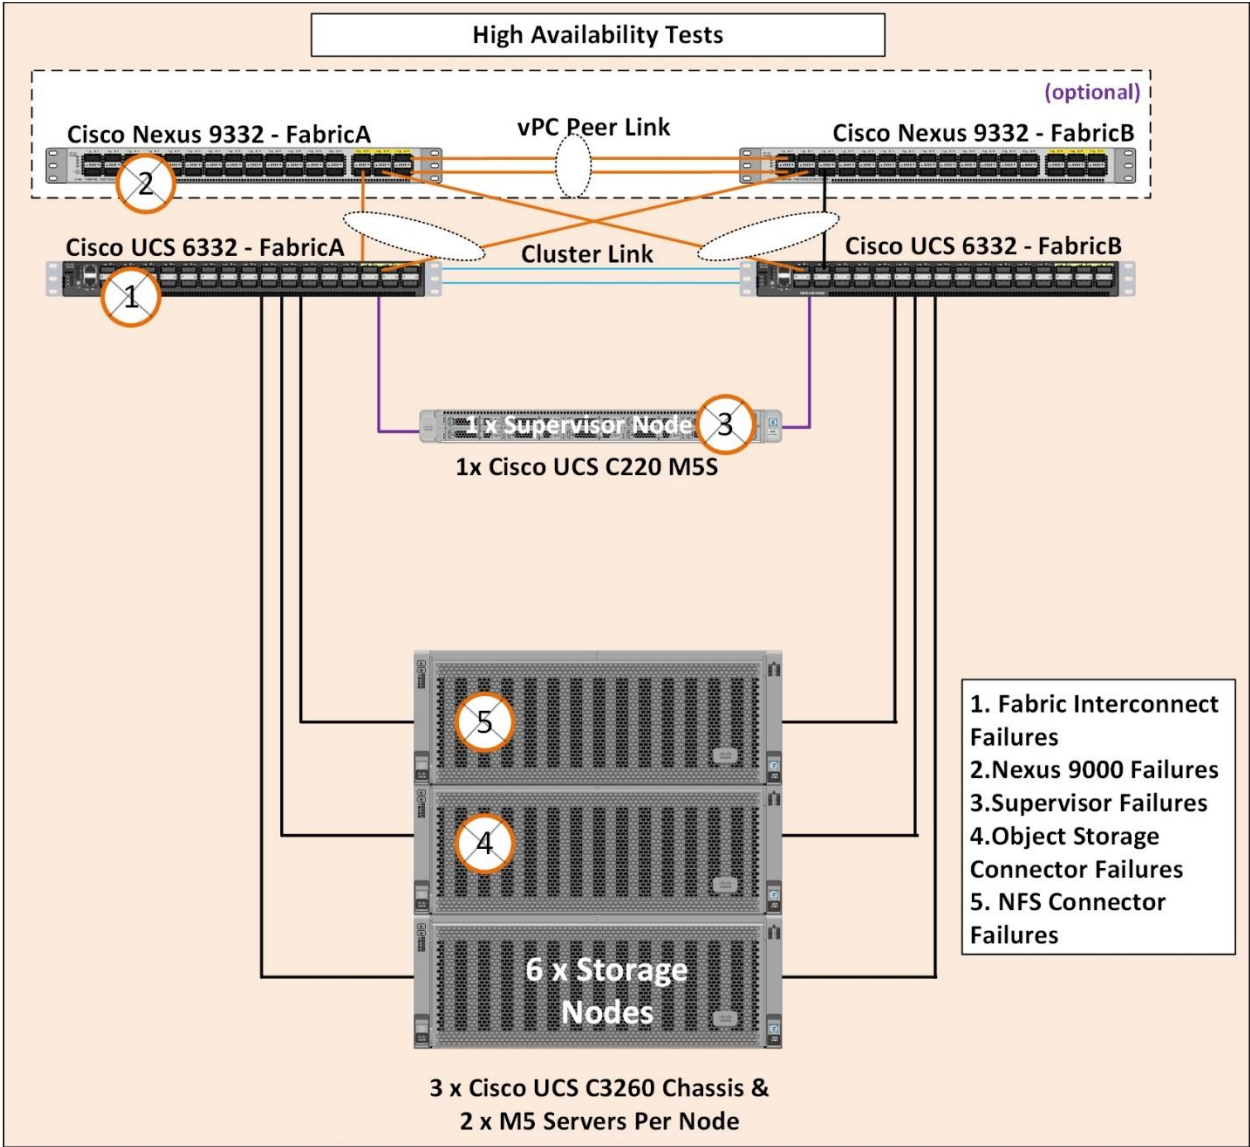
Figure 15 Network Layout



High Availability

As part of hardware and software resiliency, the following tests are in the process of being conducted on the test bed. The results of the tests will be included in the Deployment guide which will be published at a later date.

Figure 16 High Availability



Design Criteria

There are several use cases and target industries where you can use Cisco UCS and Scality's SDS RING solution. The use cases and industries are many, but not limited to the following:

Table 3 Use Cases

Primary	Backup and Archive
	Private and 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, and Cable Operator
	SaaS and Other Cloud Services
	Financial Services
	Media and Entertainment
	Police and Intelligence Agencies
	Hospitals and Medical Imaging Vendor
	Transportation
	Other Global 2000 (non XaaS, FIN, M&E, Transp, Hosp)



The following sections describe some points that may be considered for the design of the Infrastructure and the Scality RING.

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, etc.

Physical Infrastructure Considerations

Single Node versus Dual Node Cisco UCS S3260

Cisco UCS S3260 is offered with both a single node and dual node configurations for a full chassis in a 4RU rack space.

Clusters may be categorized into capacity based or performance based depending on the requirements. A dual server offers double the CPU and memory for the same set of disks; where performance is more important, a dual node configuration is recommended.

When performance is not that important and fewer cores per disk suffice as in backup or archives, a single node configuration is recommended. This will reduce the TCO of the solution.

Replication versus Erasure Coding

Scality offers the best of both worlds; replication and erasure coding, dynamic vs the size of things. The usable to raw ratio is determined by the data protection chosen. There are two standard protection models for a single site platform: Standard and Enhanced. The overhead is based on the distribution of small to large files (above and below 60KB). The more small files, the higher the overhead. The standard numbers based on field feedback are as follows:

- Standard durability has a typical overhead of 41 percent
- Enhanced durability has a typical overhead of 63 percent

Flash Storage

Flash Storage with SAS SSD's or NVMe'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 16 HDD. When using NVMe, one PCIe card is sufficient for all the drives in the S3260 systems.

As an example, a dual node configuration has 28 disk slots for each node, therefore **2 SSD's are recommended**. Note that using **2 SSD's will reduce the number of slots available per node in Cisco UCS S3260 to 26**. Using 1x2TB NVMe as mentioned in the BOM above will provide the additional 2 slots when **compared with SSD's**, this will slightly increase TB/RU of the servers.

JBOD versus RAID0 Disks

While Scality as a SDS solution works with either JBODs or RAID0 disks, it is recommended to use RAID0 for the solution. The 12G SAS RAID controller in S3260 can provide up to 4G of cache which can be used for writes. Each disk on the system is setup in one-disk RAID0 array, just so that it can benefit from the RAID Controller cache.

Memory Sizing

Memory sizing is based on the number of objects stored on each storage server, which is related to the average file size and the data protection scheme. Standard designs call for 384GB for the S3260 M5 single node, and 192GB for the S3260 M5 dual node configurations.

Network Considerations

Scality Network requirements are standard Ethernet only; refer to the Network layout in Figure 15. While Scality software can work on a single network interface, it is recommended to carve out different virtual interfaces in Cisco UCS and segregate them. A management network and Cluster network are required for the operation of the RING. Cisco UCS S3260 has two physical ports of 40G each and the VIC allows you to carve out many Virtual interfaces on each physical port.

Cisco UCS S3260 comes with 2 x 40Gb physical ports where multiple Virtual Interfaces can be configured. It is recommended to have Storage Cluster network on one port and Storage Management network on another port. The Storage Management bandwidth requirements are minimal and hence this port can be shared with other networks like Client/Application network. This will provide 40Gb bandwidth for each of these networks.

Network Uplinks

The Network uplinks from Fabric Interconnects provide upstream connectivity to the Nexus switches. A reboot for instance may be needed during a non-disruptive firmware upgrade. While there is a complete high availability built in the infrastructure, the performance may drop depending on the uplink connections from each FI to the Nexus switches. **In case you want ‘no’ or a ‘minimal drop’**, in the case of such failures or reboots, increase the uplink connections as well.

Multi-Site Deployments

Multisite deployments are available in three reference architectures from Scality:

- Stretched 2 site with Witness
- Stretched 3 sites
- 2 Single Sites asynchronously replicated

Designing a multisite is beyond the scope of this document and, for simplicity only a single site deployment test bed is setup. Please contact Cisco and Scality in case you have a multisite requirements.

Connectors

Connectors both for S3 and SOFS are installed on the storage servers, collocated with the storage processes themselves. As the RING expands for more capacity, the connector layer also expands. Most customers workload requirements are supported by this configuration.



If a customer’s workload and use case requirements not fit the assumptions made while building these standard configurations, Cisco and Scality can work together on building a custom hardware sizing to support the customer’s workload.

Expansion of the Cluster

Cisco UCS hardware along with Scality RING offer exceptional flexibility in order to scale-out as your requirements change.

Cisco UCS 6332 Fabric Interconnects have 32 ports each. Each server is connected to either of the FI's. Leaving the Network uplinks and any other servers connected to the Cisco UCS Fabric interconnects, 24 Storage server nodes can be easily configured with Cisco UCS FI pairs. If more servers are needed you should plan for a multi-domain system.

Cisco UCS offers server (KVM) management both in-band and out-of-band. If out-of-band management is planned, you may have to increase/reserve as many free OOB IP's **as needed for** managing the servers. This planning while designing the cluster makes expansion very simple.

Cisco UCS provides IP pool management, MAC pool management along with policies that can be defined once for the UCS domain. Any future expansion for adding nodes etc., 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. Once a template is created a new Service Profile for the additional server can be created and easily applied on the newly added hardware. Cisco UCS makes Infrastructure readiness, extremely simple, for any newly added storage nodes. Rack the chassis and/or nodes, connect the cables and then clone and apply the Service Profile is what needed.

Once the nodes are ready, you may have to follow the node addition procedure per Scality documentation.

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

Deployment

The test bed was deployed with 6 x Cisco UCS S3260 storage nodes with both S3 and NFS connectors. However, the deployment, performance, high availability and sizing guidelines are being worked out while this Design guide was written. More information about the deployment steps with any other best practices discovered as part of the setup will be documented in the Deployment guide.

The following is the list of activities done on the nodes while configuring the RING:

CPU Side Channel Vulnerabilities

It was recently disclosed that 3 vulnerabilities could exist that can take advantage of speculative execution in many microprocessors. The first 2 vulnerabilities are called 'Spectre', while the third as 'Meltdown'. For more information about these vulnerabilities, please click this [link](#).



Because of these vulnerabilities, the BIOS was patched and Red Hat Kernel was updated. The best practice is to have these vulnerabilities plugged.

From the test bed:

```
[root@storage-node1 ~]# uname -a
Linux storage-node1 3.10.0-693.21.1.el7.x86_64 #1 SMP Fri Feb 23
18:54:16 UTC 2018 x86_64 x86_64 x86_64 GNU/Linux
[root@storage-node1 ~]# ./spectre_meltdown.sh
```

This script is primarily designed to detect Spectre / Meltdown on supported

Red Hat Enterprise Linux systems and kernel packages.

Result may be inaccurate for other RPM based systems.

Detected CPU vendor: Intel

Running kernel: 3.10.0-693.21.1.el7.x86_64

Variant #1 (Spectre): Mitigated

CVE-2017-5753 - speculative execution bounds-check bypass

- Kernel with mitigation patches: OK

Variant #2 (Spectre): Mitigated

CVE-2017-5715 - speculative execution branch target injection

- Kernel with mitigation patches: OK
- HW support / updated microcode: YES
- IBRS: Not disabled on kernel commandline
- IBPB: Not disabled on kernel commandline

Variant #3 (Meltdown): Mitigated

CVE-2017-5754 - speculative execution permission faults handling

- Kernel with mitigation patches: OK
- PTI: Not disabled on kernel commandline

By default, Cisco UCS and Cisco UCS Manager **provide the failover capabilities to the NIC's** by connecting each node to either of the Fabrics. The disks were configured in RAID0, with battery backup (BBU) on Cisco 12G RAID controller. Cisco UCS S3260 also has redundant power supplies.

The Linux Kernel was configured with pre-requisites per Scality documentation, for example disabling Transparent Huge pages, SELINUX, etc.

Platform descriptor file from Scality simplifies installing the RING. The descriptor file is a CSV file with all the hardware information that a Scality Installer reads and automates the install. Platform CSV file was used.

The Storage nodes were also configured running the S3 and NFS connectors.

Summary

Cisco UCS Infrastructure for Scality Software-Defined Storage is an integrated solution for deploying Scality RING and combines the value of Intel Xeon architecture, Cisco data center hardware and software, along with Red Hat Linux. The solution is validated and supported by Cisco and Scality, to increase the speed of deployment and reduce the risk of scaling from proof-of-concept to full enterprise production.

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.

The Scality RING is a software-defined storage solution for petabyte-scale data storage that is designed to interoperate in the modern Software Defined Data Center. The RING software is designed to create unbounded scale-out storage systems that converge the storage of Petabyte scale data from multiple applications and use-cases, including both object and file based applications. The RING has no single points of failure, and requires no downtime during any upgrades, scaling, planned maintenance or unplanned system operations with self-healing capabilities.

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

About the Authors

Vijay Durairaj, Technical Marketing Engineer in Cisco UCS and Data Center Solutions Group, Cisco Systems, Inc.

Vijay has over 14 years of experience in IT Infrastructure, Server Virtualization, and Cloud Computing. His current role includes building cloud computing solutions, software defined storage solutions, and performance benchmarking on Cisco UCS platforms. Vijay also holds Cisco Unified Computing Design Certification.

Christopher Donohoe, Scality

Christopher Donohoe is Scality's Partner Integration Engineer. Christopher is a liaison between Scality's engineering community and the technical resources of partners like Cisco, implementing and testing new solutions prior to general availability. Christopher performs the hands-on work in documents like CVDs, while designing and architecting new automated processes for performance benchmarking and new feature validation.

Acknowledgements

- RamaKrishna Nishtala, Cisco Systems, Inc.
- Jawwad Memon, Cisco Systems, Inc.
- Tevor Benson, Scality
- Marc Villemade, Scality