White Paper

# Cisco UCS Reference Design for Object Storage

**January 2017**

# Contents

## Executive Summary

This document describes a typical object storage solution using the Cisco Unified Computing System™ (Cisco UCS®) platform. The design provides optimal deployment of software-defined object storage and networking resources all managed by Cisco UCS Manager.

The design presented in this document can be deployed with almost any object storage solution. It demonstrates how organizations can simplify object storage solutions for a variety of use cases, such as storage for unstructured data in cloud, backup, and archive workloads, by deploying Cisco UCS S3260 Storage Servers, C220 M4S Rack Servers, and 6300 Series Fabric Interconnects and Cisco UCS Manager. With Cisco UCS Manager in the data center managing the whole object storage hardware stack, you can manage all operations of the complete storage infrastructure entirely from a single pane, with industry-leading performance and scalability. Such a solution dramatically reduces both the capital expenditures (CapEx) and operating expenses (OpEx) needed to manage an object storage solution.

## Introduction

Modern data centers increasingly rely on a variety of architectures for storage. Whereas in the past organizations focused on block and file storage, today organizations are focusing on object storage, for several reasons:

- Object storage offers unlimited scalability and simple management.
- Because of the low cost per gigabyte, object storage is well suited for large-capacity needs, and therefore for use cases such as archive, backup, and cloud operations.
- Object storage allows the use of custom metadata for objects.

Table 1 summarizes the differences between traditional and object storage architectures.

**Table 1.** Comparison of Traditional and Object Storage Architectures

| | File | Block | Object |
|---|---|---|---|
| **Protocol** | Network File System (NFS) or Server Message Block (SMB) | Small Computer System Interface over IP (iSCSI) or Fibre Channel | Representational state transfer (REST) and Simple Object Access Protocol (SOAP) over HTTP |
| **Transaction Units** | Files | Blocks | Objects; files with custom metadata |
| **Metadata** | File-system attributes | System attributes | Custom metadata |
| **Manageability** | Moderate | Complex | Simple |
| **Scalability** | Unlimited | Logical unit number (LUN) size | Unlimited |
| **Best Use Case** | Shared file data | Transactional data | Archive, backup, and cloud data; so-called static data |
| **Common Storage Type** | Network-attached storage (NAS) | Storage area network (SAN) | Direct-attached storage (DAS) |
| **Cost per GB** | Midrange | High | Low |

Enterprise storage systems are designed to address business-critical requirements in the data center. But these solutions may not be optimal for use cases such as backup and archive workloads and other unstructured data, for which data latency is not especially important. However, with object storage architecture, you can easily achieve enterprise-class reliability, scale-out capacity, and lower costs with an industry-standard server solution.

The Cisco UCS S3260 Storage Server, originally designed for the data center, is now optimized for object storage solutions, making it an excellent fit for unstructured data workloads such as backup, archive, and cloud data. The S3260 delivers a complete infrastructure with exceptional scalability for computing and storage resources together with 40 Gigabit Ethernet networking. The S3260 is the platform of choice for object storage solutions because it provides more than comparable platforms:

- Proven server architecture that allows you to upgrade individual components without the need for migration
- High-bandwidth networking that meets the needs of large-scale object storage solutions
- Unified, embedded management for easy-to-scale infrastructure

Scale-out object clusters with traditional server architectures work for small businesses, but as the system grows, the complexity of management and the additional third-party network layer make this approach less compelling than Cisco UCS solutions for object storage. These traditional solutions offer no standardized integration of new hardware, and the network layer can become a bottleneck. The result is additional costs due to the need for more management or more hardware, offsetting the lower initial cost of the solution.

The Cisco UCS S3260 object storage design presented in this guide addresses today's business needs for a scalable object storage solution for unstructured data. With the power of the Cisco UCS management framework, the solution is cost effective to deploy and manage.

## Cisco UCS Solution Components

The Cisco UCS S3260 object storage design consists of a variety of components.

### Cisco UCS S3260 Storage Server

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

**Figure 1.** Cisco UCS S3260 Storage Server

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

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

The S3260 is deployed as a standalone server in both bare-metal and 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 S3260 uses a modular server architecture that, using Cisco's blade technology expertise, allows you to upgrade the computing and network nodes in the system without the need to migrate data from one system to another. It delivers:

- Dual server nodes
- Up to 36 computing cores per server node
- Up to 60 drives mixing a large form factor (LFF) with up to 28 solid-state disk (SSD) drives plus 2 SSD SATA boot drives per server node
- Up to 512 GB of memory per server node (1 TB total)
- Support for 12-Gbps serial-attached SCSI (SAS) drives
- System I/O controller with a Cisco UCS Virtual Interface Card (VIC) 1300 platform embedded chip supporting dual-port 40-Gbps connectivity
- High reliability, availability, and serviceability (RAS) features with tool-free server nodes, system I/O controller, easy-to-use latching lid, and hot-swappable and hot-pluggable components

### Cisco UCS C220 M4 Rack Server

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

**Figure 2.** Cisco UCS C220 M4 Rack Server

The enterprise-class C220 M4 server extends the capabilities of the Cisco UCS portfolio in a 1RU form factor. It incorporates the Intel Xeon processor E5-2600 v4 and v3 product family, next-generation DDR4 memory, and 12-Gbps SAS throughput, delivering significant performance and efficiency gains. The C220 M4 server delivers outstanding levels of expandability and performance in a compact 1RU package:

- Up to 24 DDR4 DIMMs for improved performance and lower power consumption
- Up to 8 small form-factor (SFF) drives or up to 4 LFF drives
- Support for a 12-Gbps SAS module RAID controller in a dedicated slot, leaving the remaining two PCI Express (PCIe) Generation 3 (Gen 3.0) slots available for other expansion cards
- A modular LAN-on-motherboard (mLOM) slot that can be used to install a Cisco UCS VIC or third-party network interface card (NIC) without consuming a PCIe slot
- Two embedded 1 Gigabit Ethernet LAN-on-motherboard (LOM) ports

### Cisco UCS Virtual Interface Card 1387

The Cisco UCS Virtual Interface Card (VIC) 1387 (Figure 3) is a Cisco innovation. It provides a policy-based, stateless, agile server infrastructure for your data center. This dual-port Enhanced Quad Small Form-Factor

Pluggable (QSFP) half-height PCIe mLOM adapter is designed exclusively for Cisco UCS C-Series Rack Servers and S3260 Storage 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 NICs or host bus adapters (HBAs). In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

**Figure 3.** Cisco UCS Virtual Interface Card 1387

The 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.

### Cisco UCS 6300 Series Fabric Interconnects

The Cisco UCS 6300 Series Fabric Interconnects are a core part of the Cisco UCS platform, providing both network connectivity and management capabilities for the system (Figure 4). The 6300 Series offers line-rate, low-latency, lossless 10 and 40 Gigabit Ethernet, FCoE, and Fibre Channel functions.

**Figure 4.** Cisco UCS 6300 Series Fabric Interconnect

The 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 6300 Series provides both LAN and SAN connectivity for all servers within the Cisco UCS domain.

From a networking perspective, the 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 NICs, HBAs, cables, and switches can be consolidated.

The Cisco UCS 6332 Fabric Interconnect is a 1RU Gigabit Ethernet and FCoE switch offering throughput of up to 2.56 Tbps 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 Small Form-Factor Pluggable (SFP) interface, these ports on a 6300 Series Fabric Interconnect can connect to another fabric interconnect that has four 10 Gigabit Ethernet SFP ports. The breakout feature can be configured on ports 1 to 12 and ports 15 to 26 on the 6332UP Fabric Interconnect. Ports 17 to 34 on the 6332 16-UP Fabric Interconnect support the breakout feature.

### Cisco UCS Manager

Cisco UCS Manager (Figure 5) provides unified, embedded management of all software and hardware components of Cisco UCS across multiple chassis and rack servers and thousands of virtual machines. It supports all Cisco UCS product models, including Cisco UCS B-Series Blade Servers, C-Series Rack Servers, S-Series Storage Servers, and M-Series Modular 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 5.** Cisco UCS Manager

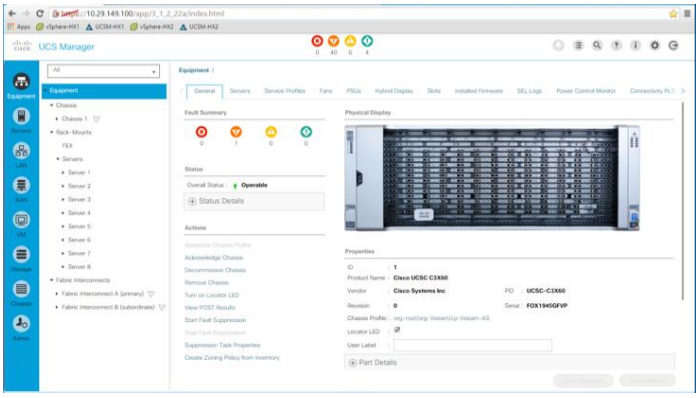


An instance of Cisco UCS Manager with all Cisco UCS components managed by it forms a Cisco UCS domain, which can include up to 160 servers. In addition to provisioning Cisco UCS resources, this infrastructure management software provides a model-based foundation for simplifying the day-to-day processes 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 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 multidomain Cisco UCS environment, enabling centralized management of distributed systems scaling to thousands of servers. The 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 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 and the Python Software Development Kit (SDK) help automate and manage configurations within Cisco UCS Manager.

## Architectural Design Considerations

Before designing an object storage solution, you need to understand how object storage works and how it can be implemented in your environment. The concept of object storage may require you to rethink the way that data is accessed in storage environments. Object storage takes a different approach from that used for traditional block and file storage.

The following sections provide a foundation to help you understand the elements that you need to consider in designing an object storage solution.

### Need for Object Storage

Traditional storage designs are well known and have been used for many years. But organizations are now shifting to object storage to address the following trends and needs:

- **Growth of unstructured data:** According to analysts such as IDC, the massive growth in unstructured data makes an object storage solution especially valuable.
- **Multidimensional scalability:** With the capability to scale resources, such as computing, network, and storage resources, in a variety of dimensions, object storage provides the flexibility of a server infrastructure with high-density storage.
- **High availability:** Typical object storage solutions can manage high availability by distributing content over multiple servers and locations.
- **Expansion without migration:** Object storage simplifies the addition of new hardware and automates the distribution of data. The use of erasure coding and the existence of multiple mirrored copies in an object storage solution enable easy re-creation of data. You can simply add new hardware and then remove the old hardware.
- **Replacement of traditional file storage:** Most object storage systems available today already have NFS, SMB, or iSCSI interfaces implemented. In addition, object storage allows you to use a unified name space across hundreds or even thousands of storage servers.
- **Standardization:** Almost all object storage systems today use the same object storage API: Amazon Simple Storage Service (S3). This common standard increases the compatibility of applications the use private object storage solutions.

### Use Cases

Object storage today is a good match for many uses cases. With its essentially limitless scalability and use of commodity hardware, it is well suited for archival storage. The capability to extend data durability with data protection methods such as replication and erasure coding demonstrates the matureness of the technology and its suitability for storing cold data.

But archiving is only one use case of many. Object storage vendors are creating new uses as the technology develops. With metadata being created for each object, additional use cases include the following:

- **Backup operations:** Essentially limitless scalability makes object storage well suited for backup operations, in contrast to typical purpose-built backup appliances, which have a limited amount of capacity.
- **Media storage:** Object storage is an excellent fit for large files such as video files. With near-line and long-term storage, an object storage solution can store all media data.

- **File-based storage:** Most vendors today offer interfaces for NFS and SMB. Although latency is still not as low as with NAS arrays, object storage solutions can host general file data for which latency is not a concern.
- **Web-scale applications:** Cloud-based applications such as consumer, social, and photo and video websites need an enormous amount of space for data storage. Object-based storage solutions are well suited for such applications because their scalability helps customers grow at a lower cost than with traditional storage solutions.
- **Virtualization:** With the adoption of flash-memory storage over the past several years for metadata and object data, performance has increased dramatically. Virtualized applications can benefit from flash-based storage, and new workloads are being adopted.

In general, all applications that require a large amount of capacity and proven data durability and that do not need low latency are well suited for object storage.

### Target Workloads

One of the most important design considerations is the target workload of your application. With the enhanced support for new object storage workloads over the past several years, object storage solutions are becoming more attractive, but they require solid understanding and planning. The rise of flash-memory storage over the past decade opened new possibilities, and the ongoing development of object storage systems also creates new possibilities. To choose among different object storage vendors and technologies, you need to understand the I/O profile of your application. Some object storage vendors concentrate on large-capacity applications and offer systems optimized for cost per amount of capacity. Some vendors are focusing on clusters with a high number of I/O operations per second (IOPS), with new methods for storing more transactional data. However, all target workloads for object storage can be categorized into three main areas— IOPS-intensive, throughput-intensive, and capacity-intensive workloads—with underlying use cases, as summarized in Table 2.

**Table 2.** Target Workloads with Use Cases

| IOPS intensive | Throughout intensive | Capacity intensive |
|---|---|---|
| • Analytics<br>• Big data<br>• Virtualization | • Web-scale applications<br>• Video on demand (VoD)<br>• Digital video recording<br>• Internet of Things (IoT) | • Archive<br>• Backup<br>• File sync-and-share |

Among these workloads, throughput- and capacity-intensive workloads are the biggest opportunities for object storage vendors, and the architectural design for these workloads differs. Throughput-intensive workloads can be reproduced by adding a small amount of flash memory to the storage layer or by using high computing power in the front-end servers, whereas capacity-intensive workloads use HDDs only.

Another benefit of object storage is that you can mix different target workloads and achieve better usability and readiness for your data center. For example, your object storage target workload could include a capacity-intensive application such as backup and a throughput-intensive web-scale application, both on the same physical solution. To build such a solution, your system must meet several requirements, such as the following:

- Your system should have different layers of storage media, such as flash-memory and disk storage.

- You need a high-performance networking solution for your servers and network infrastructure.
- You need a management layer that helps in the installation and configuration of a combined solution.

In general, object storage is gaining popularity because it can handle different workloads, allowing IT organizations to more easily incorporate it.

### Network Bandwidth

With the majority of workloads being either throughput-intensive or capacity-intensive applications, the network can easily become the bottleneck of your object storage infrastructure. Today, high-density servers with up to sixty 3.5-inch drives and throughput of more than 5 GBps per server are common. 10 Gigabit Ethernet architectures can become obstacles if they are not configured correctly and can become more expensive than 40 Gigabit Ethernet architectures. With object storage becoming a solution for more general applications, and with a typical use time horizon of three to five years, you should consider ways to help ensure that your network does not limit your infrastructure.

### Data Protection

When designing object storage solutions, another important architectural consideration is data protection. Table 3 provides an overview of the basic differences between traditional RAID protection and replication and erasure coding protection for object storage.

**Table 3.** Comparison of RAID Protection and Replication and Erasure Coding Protection

| Item | RAID | Object Storage Replication | Object Storage Erasure Coding[1] |
|---|---|---|---|
| **Protection of data** | Disk | Disk, system, and site | Disk, system, and site |
| **Overhead** | Up to 100% | Typically 200% | Up to 80% |
| **Rebuild granularity** | Per disk | Per object | Per object |
| **Rebuild time** | Depends on speed and size of disk | Depends on speed of system and size of object | Depends on speed of system and size of object |
| **Performance** | Fast | Moderate | Low |
| **Cost per GB** | Moderate | High | Low |
| **Common Storage Type** | Network-attached storage (NAS) | Storage area network (SAN) | Direct-attached storage (DAS) |
| **Cost per GB** | Midrange | High | Low |

[1] Erasure coding breaks your data into fragments and stores these fragments in different locations, such as on disks, in systems, and in different geographical locations.

RAID provides an excellent mechanism for protecting your data in a single physical system against disk failures. If you expand your solution and use multiple systems, however, RAID protection will not be able to keep up with the needs of your growing architecture. Because RAID offers only disk protection, you have to consider additional mechanisms to protect your data from a local disaster. Object storage can protect your data locally against a physical disk or system failure, and it can also protect it against a site failure. With two options—to replicate your data and to use erasure coding—object storage can help protect against both local and remote failures.

Another advantage of object storage data protection is the capability to reduce the rebuild time. With object storage, you need to rebuild only an object, whereas with RAID you need to rebuild a complete disk. With disk drive capacities reaching 10 TB and more, rebuilding a RAID group takes a long time, and the chances for bit errors increase massively. The use of replication and erasure coding over multiple disks, systems, and sites can help reduce the amount of time needed to recover data.

Keep in mind that data protection for object storage solutions can change the number of storage nodes needed in your solution, and the cluster can grow very fast. You should have a solid networking and management infrastructure that grows in parallel.

### Storage Media

As discussed in the preceding sections, object storage supports a variety of use cases and workloads, and some of these workloads use a mix of flash and disk storage. Workloads, use cases, and storage media must be suited for each other, and the use of different storage media can expand your use of object storage for applications. Not all object storage systems require different storage media, but most object storage solutions today are built on two types of storage media:

- **Flash storage:** Flash memory is predominant in traditional storage array architectures. Object storage can benefit by storing metadata in flash memory. Some object storage solutions use two subtypes of flash memory:
  - **SAS and SATA flash memory:** Dramatically reduces the latency and improves the IOPS rate and overall throughput of your object storage solution compared to general disk-based solutions
  - **Non-Volatile Memory Express (NVMe) flash memory:** Provides two or three times better performance as a result of the closer CPU architecture and direct connectivity to the PCIe bus
- **Disk storage:** When capacity and costs are more important, traditional disk drives are the best choice. With disk capacities of up to 10 TB, you can easily achieve a capacity of up to 600 TB with a high-density server.

A server such as the Cisco UCS S3260, which can work with all kinds of storage media, is clearly preferred over other solutions because it gives you more room to host different target workloads in your object storage solution.

## Physical Solution Design

The design in Figure 6 shows six Cisco UCS S3260 servers together with four Cisco UCS C220 M4S servers connected to a pair of Cisco UCS 6332 fabric interconnects. Each S3260 chassis hosts dual nodes, each with two 40 Gigabit Ethernet connections. A single node from the S3260 is connected to each 6332 fabric interconnect with a single 40 Gigabit Ethernet line. Each of the four C220 M4S servers is connected to a single 40 Gigabit Ethernet line to each 6332 fabric interconnect. The whole configuration is based on a 40 Gigabit Ethernet architecture.

**Figure 6.** Physical Solution Design for an Object Storage Solution with Cisco UCS S3260 Storage Server, C220 M4S Rack Server, and 6332 Fabric Interconnect

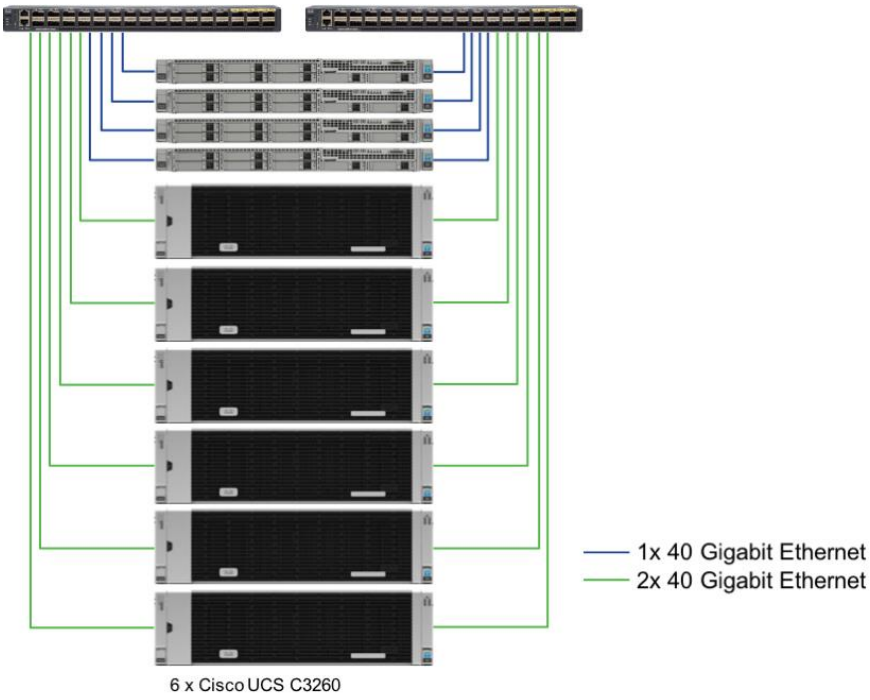


The S3260 servers are responsible for hosting the data of the object storage solution, and the C220 M4S servers are function as front-end servers. All object storage solutions have an architectural design with front-end servers acting as metadata, accessor, and gateway servers. In addition, most object storage solutions require a separate administration server to manage the solution. In the solution design shown here, three of the C220 M4S servers could be used as metadata, accessor, and gateway servers, and one C220 M4S server could be used as an administration server.

This solution design can be used for various object storage systems from different vendors. It allows flexibility where needed, but it remains still an enterprise solution.

- **Single and dual storage nodes:** The S3260 server allows single- and dual-node configurations and can handle a wide range of object storage solutions. Some solutions start with only six servers, and others start with more than six servers. The S3260 gives you flexibility either way, and it allows you to reduce the failure domain by adding a second node to the chassis.

- **HDDs:** A large number of configurations start with a minimal disk configuration. With the S3260, you can start small but still have enough room to grow to meet future needs. A general recommendation is to start with 14 disks per node.
- **Flash memory:** Business can change, and with that your requirements for object storage can change too. You may move from capacity-intensive applications to throughput-intensive or IOPS-intensive applications. With SATA, SAS, and NVMe flash storage options, the S3260 helps you manage your changing your needs. As discussed earlier, various object storage systems need various amounts of flash storage. The S3260 gives you the option to work with any object storage solution, either with or without flash storage.
- **Composable infrastructure:** The solution design presented here gives you a huge advantage when you want to incorporate new technologies or change your system. The S3260, with its composable architecture, allows you to replace individual components such as server nodes and network interfaces without the need to migrate your data. You can easily switch from an old server node to a new server node simply by replacing the node or network. You do not need to spend a long time planning for migration. This approach also reduces the overall cost of changing to new architecture generation.
- **Network:** Cisco UCS Manager and the VIC 1387 give you great flexibility in your network design. You can build a variety of virtual NICs (vNICs) for different VLANs in your object storage solutions, and with additional configuration of each vNIC for QoS, class of service (CoS), and maximum transmission unit (MTU) size, you can individually size and manage your solution. For example, your administration network may need an MTU size of 1500, a QoS level of Bronze, and CoS of 0; your cluster network with S3260 servers may have an MTU size of 9000, a QoS level of Platinum, and a CoS of 3.

## Target Workload Designs

Using the physical design discussed in the previous sections as a basis, you can also create several logical designs using the Cisco UCS S3260 for an object storage solution based on target workloads. The main advantages of the S3260 are its flexibility in hosting various object storage workloads, its high network performance, and its simplicity of management. The Cisco C220 M4S front-end rack servers can be considered fixed-configuration servers, whereas the S3260 servers can be customized in a variety of ways to meet different target workloads.

### Logical Design for an IOPS-Intensive Object Storage Solution Workload

The logical design example in Figure 7 shows how the S3260 could be configured for an IOPS-intensive object storage solution. Each node in a chassis has a relative high amount of flash storage to achieve a higher IOPS rate. The remaining slots are configured with disk drives to store the data. Note that the ratio of SSDs to HDDs can vary depending on the IOPS rate that you want to achieve with your solution, but an overall SSD:HDD ratio of 1:2 is a common practice.

**Figure 7.** Logical Design Example for Cisco UCS S3260 for an IOPS-Intensive Object Storage Solution


## Logical Design for a Throughput-Intensive Object Storage Solution Workload

Figure 8 shows an example of a throughput-intensive object storage solution. The configuration of the S3260 is slightly different than for the IOPS-intensive solution. The SSD:HDD ratio is a more conservative 1:4 because the overall solution focuses on increased throughput while retaining the capability to handle some amount of transactional data.

**Figure 8.** Logical Design Example for Cisco UCS S3260 for a Throughput-Intensive Object Storage Solution

## Logical Design for a Capacity-Intensive Object Storage Solution Workload

For the capacity-intensive object storage solution in Figure 9, the S3260 is configured in HDD-only mode, hosting 56 disks with 28 disks per node. This configuration provides the maximum raw capacity, with up to 560 TB per chassis and with an overall raw capacity of 3.4 PB for the entire reference design configuration. If you choose erasure coding for data protection, you can easily get over 2 PB of usable capacity.

**Figure 9.** Logical Design Example for Cisco UCS S3260 for a Capacity-Intensive Object Storage Solution

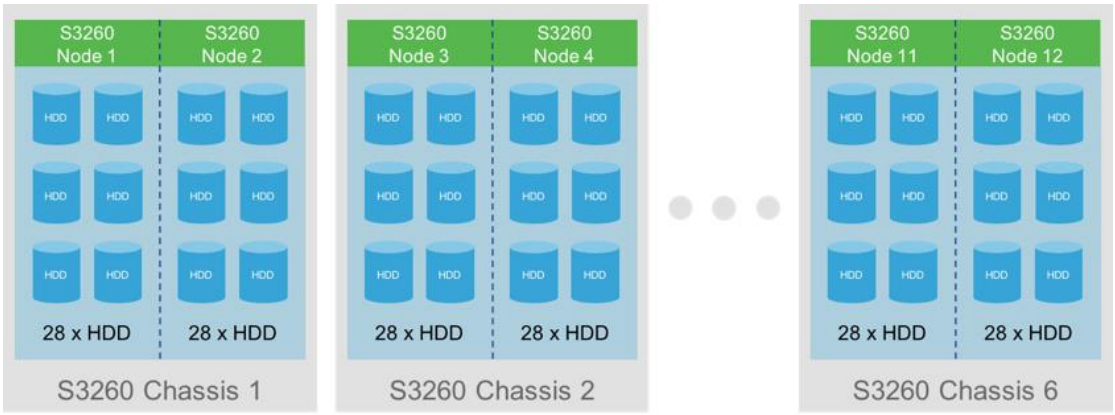


## Logical Design for a Mixed-Workload Object Storage Solution

Figure 10 shows an example of a mixed-workload object storage solution configured on an S3260. By using different configurations for different target workloads, you can easily map a variety of business requirements on the S3260 solution. The example here shows two workload designs on the same solution. One design supports a throughput-intensive workload with a mix of SSDs and HDDs. The other workload uses a pure disk design.

**Figure 10.** Logical Design Example for Cisco UCS S3260 for a Mixed-Workload Object Storage Solution

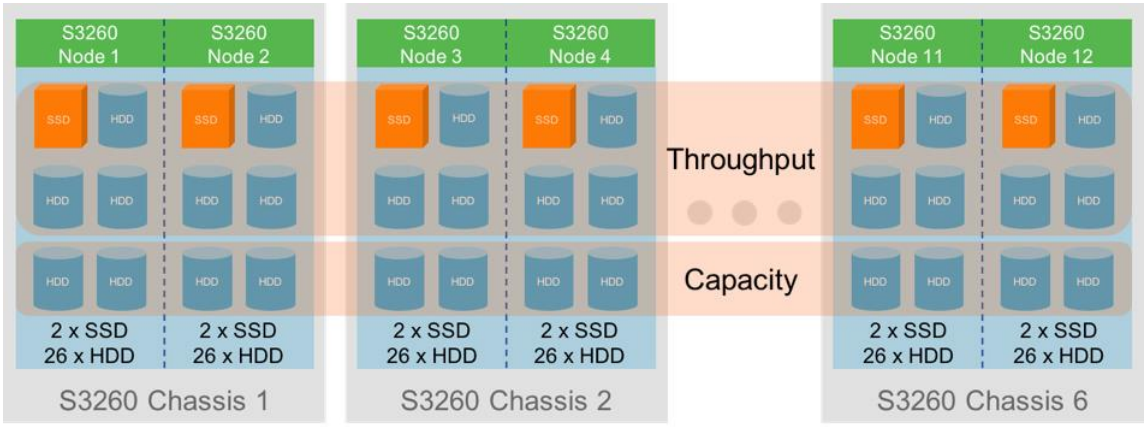


The S3260 provides options to support all the different workloads, letting you mix disk technologies such as NVMe SSDs, SAS SSDs, and near-line SAS (NL-SAS) HDDs with a 40 Gigabit Ethernet network. It also offers the benefits Cisco UCS Manager, which lets you easily add nodes to meet your business needs.

## Conclusion

With the rapid growth of unstructured data such as archive and backup data, traditional storage arrays are no longer meeting the need for efficient high-capacity storage. Organizations thus are now turning to object storage solutions. Object storage, with the advantages of easy scalability and cost-effective petabytes of storage at scale, is an increasingly mature technology. With its additional development in enterprise management, object storage no longer needs to hide its open-source background.

The Cisco UCS platform offers an excellent solution for the high-demand operations of object storage. It provides a low-cost, flexible, enterprise management solution with the reliability needed to store unstructured terabytes and petabytes of data at scale. The Cisco UCS S3260 Storage Server, together with the Cisco C220 M4S Rack Server, Cisco UCS 6300 Series Fabric Interconnects, and Cisco UCS Manager, offers an excellent self-sufficient infrastructure for object storage solutions. You can design a highly flexible solution for almost any object storage solution on the market. The flexibility of the S3260 offers you a wide variety of ways to meet your business needs and to help ensure that you get an optimal solution for your target workload.

With a common reference design for object storage, the Cisco UCS solution offers the simplicity and the flexibility you need to support a variety of workloads and requirements.

## For More Information

For additional information, see the following:

- Cisco UCS S3260 Storage Server: http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-s-series-storage-servers/index.html
- Cisco UCS C220 M4 Rack Server: http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c220-m4-rack-server/index.html
- Cisco UCS 6300 Series Fabric Interconnects: http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html
- Cisco UCS Virtual Interface Card 1387: http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1387/index.html
- Cisco UCS Manager: http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**



Printed in USA C11-738251-00 01/17