

Cisco Data Intelligence Platform with MinIO



Existing data landscape and architectural evolution

Data scientists are constantly searching for newer techniques and methodologies that can unlock the value of big data and distill this data further to identify additional insights that can transform productivity and provide business differentiation. This complex effort covers everything from storage, security, resilience, access, analytics, and APIs, in addition to operational considerations. The goal is to create an environment in which the business can transform itself into a data-driven enterprise.

At the heart of this data-driven journey is Artificial Intelligence and Machine Learning (AI/ML). AI/ML is a transformative technology that demands performance and flexibility in the form of different types of compute (CPU, GPU, and FPGA) coupled with modern, scalable, and performant object storage. This new architecture stands in stark contrast to a traditional data lake. The inability to process, transform, and perform analytics in a single place has led to a balkanization of data. As data grows, the problem grows with it, creating silo after silo. The impact on an enterprise is profound, as internal teams battle over what represents the truth—creating massive analytical and operational inefficiencies.

These challenges are compounded by the rapid increase in unstructured and semi-structured data. Given the volume, these data types demand cost-effective and performant approaches to storage and management. Through data-protection schemes such as erasure coding and bit rot protection, object storage has become the primary choice for enterprises seeking to store and access massive amounts of data in an economical fashion.

Highlights

Cloud-scale architecture unifies advanced analytics, big data, and Artificial Intelligence (AI) workloads using modern, high-performance object storage

Cisco and MinIO bring together diverse, siloed workloads under a scalable architecture that delivers a robust experience at both the infrastructure and platform levels.

Paradigm shift enables diverse computing constructs to work on data

The architecture of the Cisco® Data Intelligence Platform solution enables data to be operated by different computing constructs—whether CPU or Graphics Processing Unit (GPU) or Field-Programmable Gate Array (FPGA)—based on application demand.

The Cisco Data Intelligence Platform unifies these silos by creating an architecture designed to deliver performance at scale. This iteration of the highly successful architecture pairs the Data Intelligence Platform with high-performance object storage from MinIO and Cloudera's Data Science Workbench.

With the addition of modern object storage from MinIO, data architects and IT can scale storage and compute seamlessly and infinitely, effectively manage the lifecycle of data, offer the computational flexibility to handle a range of AI/ML workloads, and deliver superior total cost of ownership.

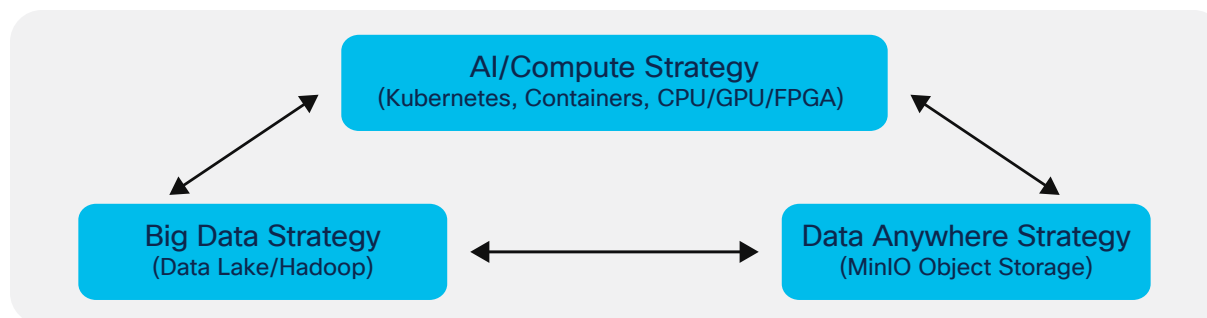
Cisco Data Intelligence Platform

The Cisco Data Intelligence Platform brings big data, AI compute farms, and tiered storage to work together as a single entity, but with each element able to be scaled independently. This architecture supports:

- Extremely fast ingestion and engineering of data performed at the data lake
- An AI computing farm, allowing different types of AI frameworks and computing resources (GPU, CPU, and FPGA) to work on this data for additional analytics processing
- A storage tier, allowing the gradual retirement of data that has been worked on to a dense storage system with a lower cost per terabyte, reducing TCO

The Cisco Data Intelligence Platform supports today's evolving architecture (Figure 1). It combines a scalable infrastructure, centralized management, and a fully supported software stack (in partnership with industry leaders in relevant areas) to each of these three independently scalable components of the architecture: the data lake, AI/ML technologies, and object stores. (See Figure 2.)

Figure 1. Cisco Data Intelligence Platform architecture



Independently scale storage and computing resources based on demand

The combined architecture enables both data-intensive and computation-intensive workloads to work together. Customers can start small and expand non-disruptively, independently scaling storage and/or computing resources to hundreds of thousands of nodes.

MinIO Object Storage

MinIO is a high-performance, distributed, software-defined object storage system. It is open source under the Apache v2 license. MinIO is purpose-built to server-only objects, currently making it the fastest private cloud object store.

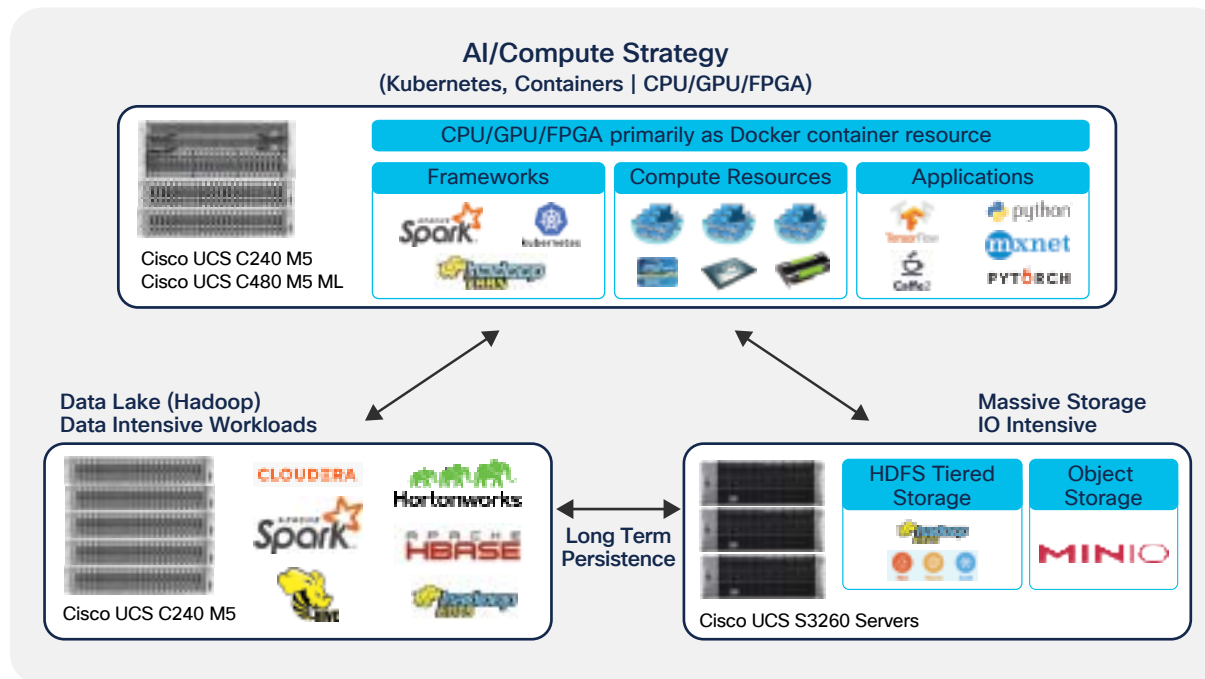
MinIO bit rot protection

MinIO's advanced capabilities in erasure code and bitrot detection mean that an enterprise can lose up to half the servers and continue to serve data. Data protection code is accelerated using SIMD instructions on x64 CPUs.

MinIO

MinIO's object storage server augments Cisco's Data Intelligence Platform architecture by providing exceptional throughput and scale, effectively creating a second tier of warm data to complement the Hadoop tier. With MinIO object storage, on-premises enterprise customers can create highly performant, cost-effective solutions to support a variety of AI/ML workloads and use cases. For example, new data tends to be read more frequently than older data. New data is deemed hot, and older data is considered warm. With the advent of high-throughput object-storage systems, more data is accessed more frequently, creating a massive tier of warm data. While hot data receives the bulk of the attention in most enterprises, it is warm data that offers the most potential for insight generation and pattern recognition given its combination of recency and scale. Optimizing this ever-growing tier from a performance and cost perspective is critically important.

Figure 2. Cisco Data Intelligence Platform with MinIO and software stack and ISV partners



MinIO cloud-native capabilities

MinIO was built from scratch over the past four years. It was designed to run natively in containers and be managed using orchestration services such as Kubernetes, delivering simplicity and elasticity for dynamic workloads.

Architecture built on Cisco UCS

The Cisco Data Intelligence Platform is built on Cisco UCS®, a proven platform for enterprise analytics applications. Cisco UCS offers a broad portfolio of Cisco Validated Designs with industry-leading, Independent Software Vendor (ISV) partners in each of these areas: big data, AI, and object storage.

Ease of deployment

Cisco UCS Manager, together with Cisco Intersight™, simplifies infrastructure deployment with an automated, policy-based mechanism that helps reduce configuration errors and system downtime.

In the architecture in Figure 2, each persistent store (HDFS with Cloudera and Object Store with MinIO) are capable of delivering sufficient throughput speeds to handle a wide variety of AI/ML workloads. More importantly, the advent of a high-performance object storage tier brings significant advantages in terms of cost, complexity and APIs.

The hot-data tier delivers a storage tier that consists of Cisco UCS C240 M5 Rack Servers to store data sets that require high-speed storage access. The object storage warm-data tier uses Cisco UCS S3260 M5 Storage Servers, providing high-capacity, high-speed storage. Each Cisco UCS S3260 chassis is equipped with dual server nodes and has the capability to support up to hundreds of terabytes of MinIO erasure-coded data, depending on the drive size. MinIO is optimized for large data sets used in scenarios such as big data, cloud, analytics, video surveillance, and content delivery environments.

MinIO is an S3-compatible, software-defined, distributed object storage system. It is comprised of an object storage server, a client, and SDK. MinIO's performance characteristics allow analytical frameworks like Spark, Presto, and TensorFlow to operate directly on the object store.

In addition to its industry-leading performance, MinIO has developed a rich set of enterprise features:

- MinIO's advanced capabilities in erasure code and bit rot detection mean that an enterprise can lose up to half its servers and continue to serve data. Data protection code is accelerated using SIMD instructions on x64 CPUs.
- MinIO supports sophisticated, low-overhead encryption and tamper-proofing that exceeds what is available from leading cloud providers.
- MinIO's multi-site federation supports an unlimited number of instances to form a unified global name space. Federation is often paired with continuous replication for large-scale, cross-data-center deployments. By leveraging Lambda compute notifications and object metadata, it can compute the delta efficiently and quickly, keeping data loss to a bare minimum should a failure occur—even in the face of highly dynamic data sets.

- MinIO is cloud-native, designed to run in containers and to be orchestrated with Kubernetes—delivering performance, elasticity, and API compatibility with the rest of the cloud-native stack. This facilitates the collaboration between development, operations, and IT.

Because MinIO is purpose-built to serve only objects, a single-layer architecture achieves all of the necessary functionality without compromise. The advantage of this design is an object server that is simultaneously high-performance and lightweight.

Cisco UCS S3260 Storage Server

The Cisco UCS S3260 Storage Server (Figure 3) is a modular storage server with dual server nodes. It is optimized for large data sets used in scenarios such as big data, cloud, object storage, video surveillance, and content delivery environments.

The S3260 server helps achieve the highest levels of data availability and performance. With a dual-node capability that is based on the latest second-generation Intel Xeon and Intel Xeon Scalable processors, it offers up to 840 Terabytes (TB) of local storage in a compact, 4-rack-unit form factor. Network connectivity is provided with dual-port 40-Gbps nodes in each server.

Figure 3. Cisco UCS S3260 Storage Server



Reference architecture

Tables 1 and 2 in the [Cisco Data Intelligence Platform solution brief](#) summarize the reference architecture configuration details for the data lake, AI/ML components of the data lake, and object storage.

Table 1 in this document lists the reference architecture for MinIO object storage.

Table 1. Cisco UCS Integrated Infrastructure for Big Data and Analytics configuration options

Component	High capacity
Servers	Cisco UCS S3260 Storage Servers
CPU	2 x 2nd Generation Intel Xeon Gold 6230R Processor (2 x 26 cores and 2.10GHz)
Memory	12 x 32-GB 2933 MHz DDR4 (384 GB)
Boot	2 x 240-GB SATA SSDs
Storage	28 x 8TB 7.2K rpm LFF SAS HDDs
VIC	40 Gigabit Ethernet (Cisco UCS VIC 1300)
Storage controller	Cisco UCS S3260 dual RAID controller

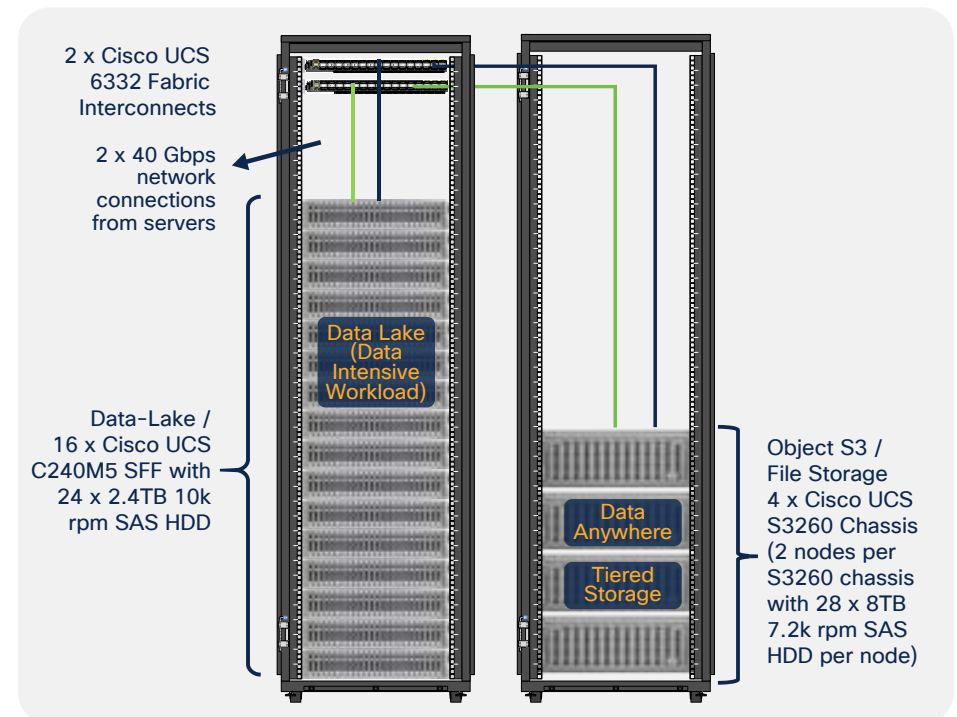
Sizing MinIO on Cisco UCS

It is recommended to configure a storage usage ratio of 1.33 for MinIO. For example, to deliver one petabyte (1 PB) of usable data, MinIO will need 1.33 PB of raw disk capacity.

As such, with four Cisco UCS S3260 chassis (eight nodes) and 8-TB drives, MinIO would provide 1.34 PB of usable space (4 multiplied by 56 multiplied by 8 TB, divided by 1.33).

Figure 4 illustrates an eight-node cluster with a rack on the left hosting four chassis of Cisco UCS S3260 M5 servers (object storage nodes) with two nodes each, and a rack on the right hosting 16 Cisco UCS C240 M5 servers (Hadoop data lake). Each link in the figure represents a 40-Gigabit-Ethernet link from each of the servers directly connected to a fabric interconnect. Every server is connected to both fabric interconnects.

Figure 4. Petabytes of Hadoop data consolidated in a single rack with 8 Cisco UCS S3260 M5 servers



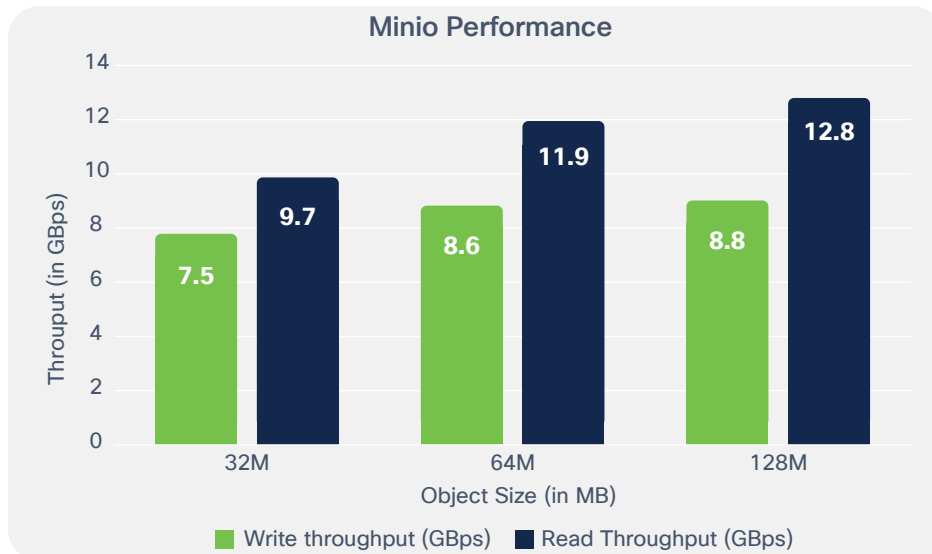
Performance benchmark

The [S3-benchmark tool](#) was used to evaluate MinIO's performance. Object storage systems are based on HTTP and the Representational State Transfer (REST) protocol. In addition, MinIO object storage systems support other important HTTP-based protocols, such as Simple Storage Service (S3) APIs, which is a standard interface for object storage. Object storage can scale to thousands of petabytes and costs less to manage.

The benchmark setup is comprised of four MinIO servers (on two Cisco UCS S3260 chassis) and four S3-benchmark clients (Cisco C240M5 servers) connected to 40-Gbps network interconnects, as shown in the topology in Figure 5, to measure the S3 benchmark performance numbers for the READ and WRITE throughput. Both client and inter-node communications provide a maximum of 40 Gbps, which equates to 5 GB per second (GBps).

Maximum aggregate throughput was seen when S3-benchmark was run with 64 threads at a 128-MB object size. The aggregate write throughput was 8.8 GBps and the aggregate read throughput was 12.8 GBps.

Figure 5. S3-benchmark performance



The performance benchmark command used was:

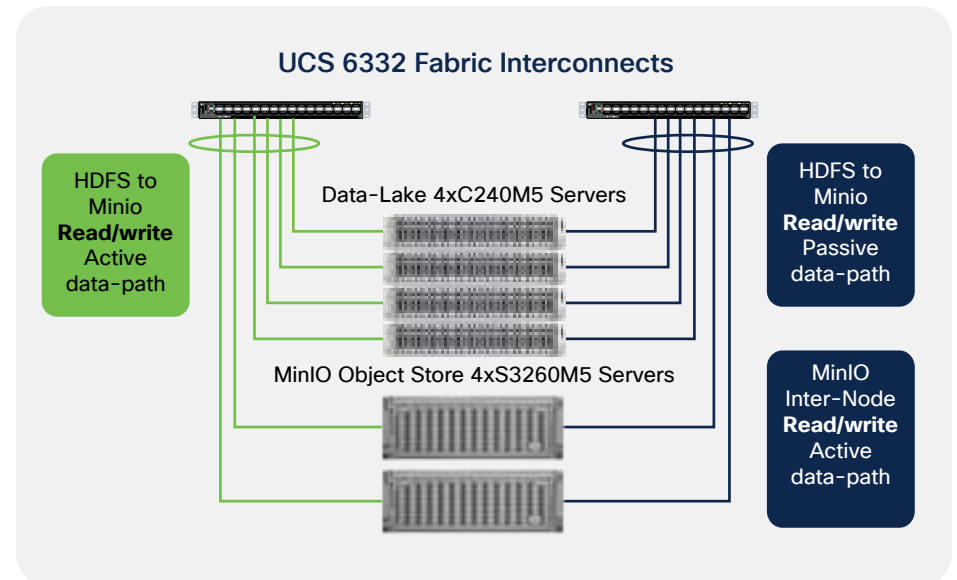
```
./s3-benchmark -u http://minio-x:9000 -b $(hostname) -z 128M -t 100 -d 30 -a minio -s minio123
```

Topology

Lab topology consists of 4x S3260 servers containing 112 total drives. Each drive was 8 TB in size and the MinIO is configured with a storage usage ratio of 1.33.

To mimic the real-time use case, drives were equally divided into seven sets of 16 drives with the storage use ratio being 1.33, while objects are sharded across 12 data and four parity blocks. In this configuration, MinIO can lose up to four of the 16 drives without losing data. Higher storage ratios would allow for up to 8 drive failures before data is lost (Figure 6).

Figure 6. MinIO S3-benchmark test topology



Conclusion

MinIO on Cisco UCS provides a proven, high-performance solution that is easy to scale and manage while providing better TCO to the customer. As drives become denser and networking becomes faster, the accompanying object storage solution must be performant, scalable, cost-effective, and easy to manage. The Cisco UCS S3260 M5 Storage Chassis provides high-capacity storage with two S3260 server nodes supporting hundreds of terabytes of inexpensive, yet performant object storage. This makes a clear case for the Hadoop stack to operationalize data tiering from the Hadoop Distributed File System (HDFS) to MinIO, offering an attractive alternative for today's critical workloads. The current design offers a proven deployment model for enterprise Hadoop while enabling a second, highly economical second tier of warm storage.

The configuration detailed in this document can be extended to clusters of various sizes, depending on the application demands. Scaling beyond a single rack (eight servers) can be implemented by interconnecting multiple Cisco UCS domains using Cisco ACI® technology. This architecture is scalable to thousands of servers and to hundreds of petabytes of storage and can be managed from a single pane using Cisco Intersight.

For more information

To learn more about Cisco Data Intelligence Platform, visit: <https://www.cisco.com/c/dam/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/solution-overview-c22-742432.pdf>.

To find out more about the MinIO Object Storage, visit: <https://min.io>.

MINIO®, the MINIO® logo design (bird) and the MINIO wordmark design are exclusive trademarks registered by MinIO, Inc.