



SEAGATE AND MINIO SOLUTION BRIEF

# The Data Density Challenge

TECHNOLOGY PAPER

## EXECUTIVE SUMMARY

The strategic imperatives facing the modern enterprise are truly daunting. Digital transformation, big data/analytics, machine learning (ML) and artificial intelligence (AI) each bring unique challenges. Each imperative demands a solution that is tailored to its unique computing, networking, and storage requirements. How then does an enterprise design a scalable infrastructure which can bridge these different technology roadmaps?

The answer lies in combining fast networking and dense storage with modern, high-performance, cloud-native object storage software. Together, these components can serve as the foundation for a wide variety of enterprise workloads, from emerging AI and ML tasks to more standard large-scale archival and backup workloads.

This paper details the key parts of this architecture and how technology from Seagate and MinIO is altering long-held beliefs about the shape of the price/performance curve.

## THE BUILDING BLOCKS OF DATA DRIVEN ENTERPRISE

The velocity of data creation continues to accelerate. In a Seagate-sponsored white paper, IDC estimated that between 2018 and 2025 data will grow more than five-fold from 33 zettabytes (ZB) to 175 ZB. For that data to be useful to the enterprise, it has to be available to the enterprise. The concept of availability is often overlooked—primarily because it is multifaceted.

Availability means the data must be protected, strictly consistent, accessible, and fast. Many solutions and technology combinations can address one, two, or occasionally three of these solutions, but Seagate and MinIO are unique in their ability to address all the attributes of availability. Together, they combine dense compute and storage with modern, high-performance object storage. The result is a solution that can scale up to *7.5PB of usable capacity per rack*, and can be made accessible over a 100Gb/s network. MinIO can run two distinct instances per each Seagate® Exos™ E 4U106 chassis. The two nodes logically split the disk array among the two servers attached to the same chassis. This is called a split-mode, shared-nothing architecture. This architecture provides exceptional flexibility and high performance for a variety of big data and AI workloads.


The result is an entirely new price/performance curve for the enterprise. Across retail, finance, healthcare, government, security, and media/entertainment enterprises, architects are using the combination of Seagate enterprise storage systems and MinIO to target a spectrum of workloads from Spark and Presto to archival and snapshot backups. The simplicity of the MinIO and its open-source Apache License V2 coupled with the superb economics and reliability of the Seagate Exos E 4U106 deliver exceptional total cost of ownership (TCO) results.

# SEAGATE

Seagate® Exos™ E 4U106 is the data center's largest building block, delivering industry-first capacity and density without sacrificing data access speed. The vast amount of data that fits in a 4U rack by leveraging up to 106 Seagate industry-leading, high-capacity, 16TB nearline hard drives in a single enclosure that holds up to an unprecedented 1.7PB of business intelligence. With an overall maximum bandwidth of 36GB/s, you can access mission-critical and archival data with lightning speed and be assured about data availability through hot-swappable controller, PSUs, fans, drives, and expander cards.



The Exos E 4U106 JBOD is suited for data storage requirements that are expected to grow rapidly. When combined with two external 1U servers connected to the Exos E 4U106 via the 12Gb/s SAS links, the Exos E 4U106 running in a split-chassis, shared-nothing mode provides two 54-drive nodes with high-performance compute for the MinIO object storage server. Innovative drive placement maximizes airflow and minimizes power consumption, and unique performance, efficiency, and scalability features provide an unprecedented low TCO. The system's new modular, user-focused design minimizes system downtime. Its modularity makes it first-in-class for reliability, easy setup, maintenance, and expansion via hot-swappable expanders and dual data paths, as well as redundant I/O modules, fans, and power supplies.



## MINIO OBJECT STORAGE

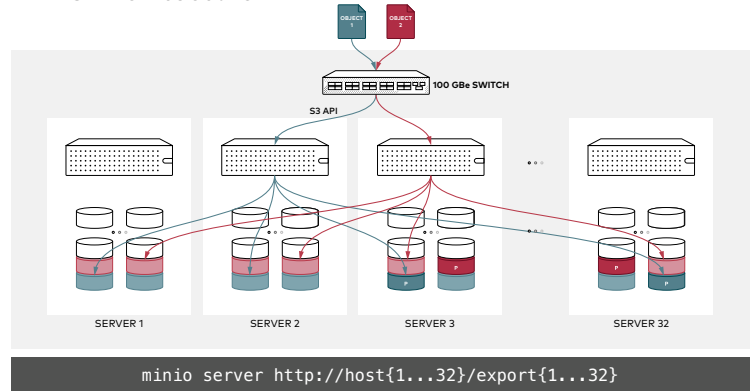
The MinIO object storage server is designed to store all types of high volume and velocity data, from time series data and log files to data for analytics workloads. MinIO's software-defined storage provides unlimited scalability through global federation and the industry's best published performance. This makes MinIO an ideal S3-compatible replacement for Hadoop HDFS for ML, AI, and other big-data workloads, as well as a performant, economical private cloud alternative to expensive public cloud instances.

MinIO was designed only to serve objects, which in turn drives its exceptional performance. MinIO was also designed for the enterprise with a suite of features that include full S3 compatibility, support for S3 Select, Encryption, WORM, Bit-rot Protection, Identity Management, Continuous Replication, Global Federation, Lambda Compute and Multi-Cloud Gateway support. In particular, MinIO's use Intel Xeon SIMD instructions in the Seagate Exos E 4U106 to enable frictionless encryption, erasure-code, bit-rot hashing, and compression.

## ARCHITECTURE AND PERFORMANCE

MinIO object storage delivers its performance by fully leveraging I/O resources. In our tests running S3 workloads, MinIO could leverage the entire bandwidth available from the drives, delivering 85MB/s sustained throughput from each hard drive. As the number of drives was increased, the combined throughput of the drives exceeded the capacity of the network, at which point MinIO saturated the network bandwidth in its entirety.

MinIO Architecture



## CONCLUSION

It is hard to predict the future—particularly when technology is involved. It helps to start with what we are most confident about. The continued growth of data is one of those building blocks. The power of AI/ML and analytics is another. Increased automation is a third. For an enterprise to succeed, it must position itself to rise to the challenges presented by all of these trends.

Together, MinIO and Seagate provide enterprises with a remarkably powerful and flexible solution to address these challenges. Their software and hardware offerings support data growth through efficient, modular products while delivering ultra-reliable object storage performance, data availability, and cloud-native functionality.

## ABOUT SEAGATE

For 40 years, Seagate has enabled exponential data growth with breakthrough hard drives, solid state drives, systems, and recovery services. We provide end-to-end data management solutions across surveillance, NAS, data centers, consumer data storage, and more. Seagate creates space for the human experience by innovating how data is stored, shared, and used. Learn more at [seagate.com](https://www.seagate.com).

Seagate developed Lyve™ Data Labs to provide a collaborative workspace to enable enterprises across industries to fast-track data transformation and solve the challenges of IT 4.0. Ready to learn, engage, and build with Seagate? [Visit Lyve Data Labs](#).

## ABOUT MINIO

Founded in 2014, MinIO is now the world's fastest-growing object storage system. Backed by some of the smartest minds in storage and venture capital, including Nexus, General Catalyst, Dell Technologies Capital, Intel Capital, AME Cloud Ventures, and key angel investors, the company has raised \$23.3M through its Series A round. Learn more about MinIO at [min.io](https://min.io).



[seagate.com](https://www.seagate.com)

© 2019 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Exos, the Exos logo, Lyve, and the Lyve logo are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Seagate reserves the right to change, without notice, product offerings or specifications. SB508.2-1911US, November 2019