

Meet Data Challenges Affordably with Seagate® Lyve® Rack and Intel® Technologies

Lyve Rack, powered by Intel® Xeon® Scalable processors, delivers a high-capacity object-storage solution that organizations like yours can easily deploy to manage increasing data growth

Intel Technologies

Intel Xeon Scalable processors provide the performance necessary for fast write, read, and search capabilities in Seagate Lyve Rack. Improvements in 2nd Generation Intel Xeon Scalable processors over previous-generation Intel Xeon processors include:

- Improved performance through higher core counts and higher maximum Intel Turbo Boost Technology frequencies^{1,2}
- Higher memory speeds³
- Larger memory capacity⁴
- More cache⁵

As the CORTX community grows, Intel® Optane™ technology will become an important part of many object-storage solutions built using CORTX.

What insights are you missing due to missing data?

Data is being created every day at a ferocious rate. In a study commissioned by Seagate in 2018, IDC estimated that the amount of data created, replicated, and consumed worldwide would grow from 33 ZB to 175 ZB by 2025. That much data is enough to fill 12.5 billion of today's largest hard drives.⁶ However, storing unstructured data has traditionally been too expensive to justify saving and using more than a small fraction of the data. While hyperscale storage services can be extremely economical for storing vast quantities of data, accessing that data can entail unexpectedly high charges. By contrast, enterprise object-storage solutions impose no extra charges to access data, but they tend to be too costly overall to store massive amounts of data. Until now, the economics of both saving and using enormous quantities of data have not aligned, and companies have been forced into unpleasant choices about which data to keep and which to discard.

This profligate waste of data represents countless missed opportunities for businesses and society. Many business applications—particularly artificial intelligence (AI) applications—perform better with more data

available for analysis, training, classification, and prediction. With so much data being disposed of without analysis, what insights are being missed by your business? What value could you uncover if you could economically retain more of your unstructured data?

CORTX Open Source Object Storage

CORTX is an open source, software-defined object store sponsored by Seagate. Seagate developed CORTX to offer low-cost object storage using the latest generation hardware for enterprise customers, allowing them to capture the efficiency and experience of the cloud at the best possible price.

A community of data scientists and big data and enterprise storage experts designed, built, and maintains CORTX. They built it with an eye toward future data center requirements: scalability, resiliency, and hardware efficiency.

Key features of CORTX include:

- A familiar, Amazon S3-compatible API (S3 protocol) for local data access
- An identity and access management (IAM)-compatible API for user management and data protection using bucket policies and access-control lists (ACLs)
- Multitenancy
- A management layer with an API, command-line interface (command-line tool), and graphical user interface (GUI)

- Integration with the Seagate Lyve ecosystem for full lifecycle data management from ingestion to movement to activation of your data
- Data protection using hierarchical erasure coding

Figure 1 shows how CORTX fits within the ecosystem of the network edge, AI, 5G, and machine-to-machine communications. The key to all of this functionality is the CORTX object-storage architecture.

CORTX Architecture Overview

CORTX features an inherently scalable architecture. Each server node within an object-storage system built on CORTX connects to the wider data center and storage infrastructure through management, network, and storage connections.

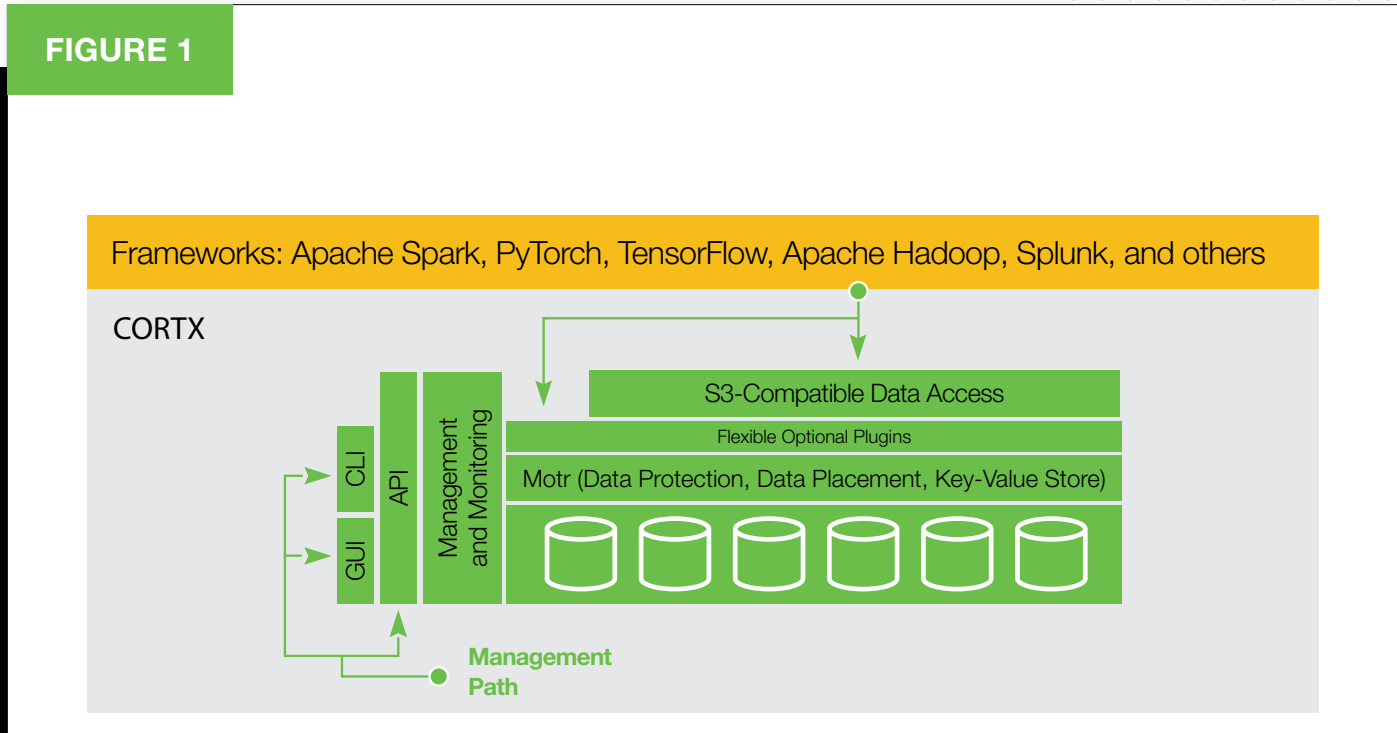


Figure 1. The role of CORTX in the application ecosystem



FIGURE 2

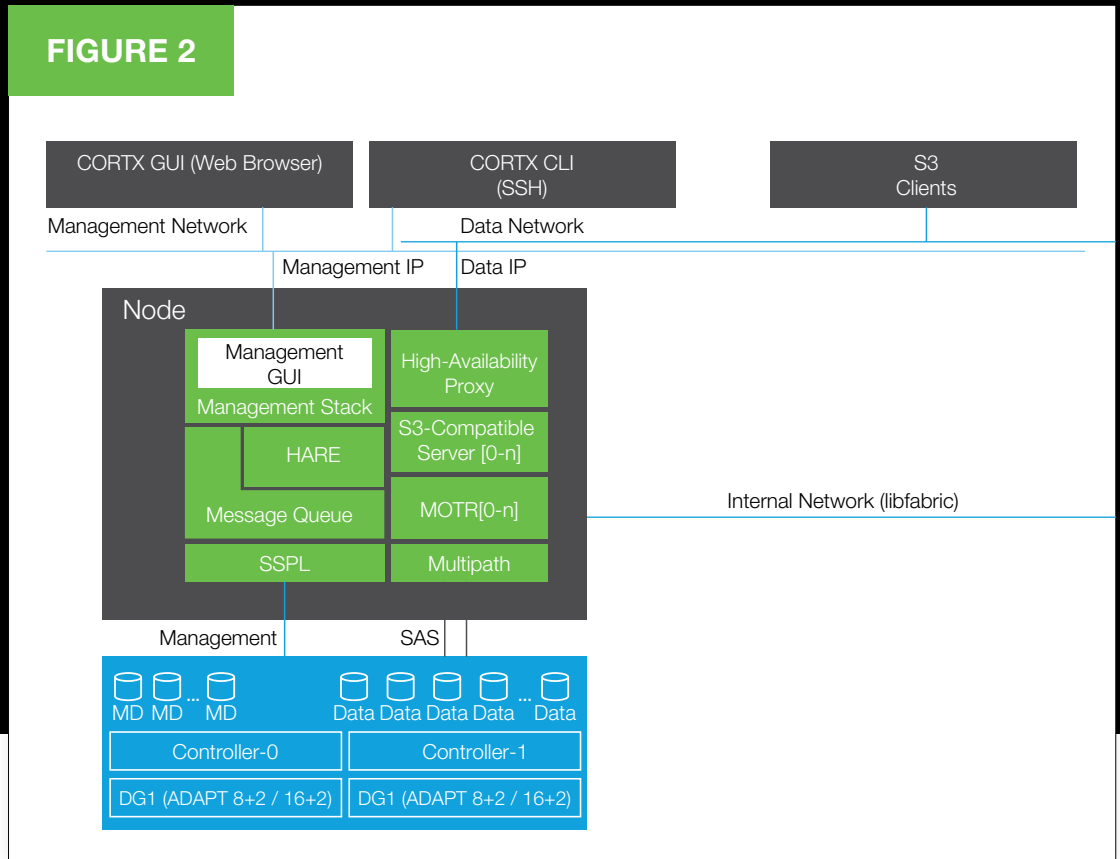


Figure 2. Architectural diagram of CORTX

Management for CORTX is performed through the Component Service Management API, GUI, or command-line interface (CLI), as shown in Figure 2. Management continues through the CORTX node to the storage backend using the (optional) Seagate Storage Platform Library (SSPL). This enables administrators to manage their block storage without having to rely on a separate software layer within the storage hardware.

HARE is the software in CORTX that configures the Motr object and key-value store. HARE starts and stops Motr services, and it is ultimately responsible for high availability for Motr and the S3 protocol server.

The actual storage in CORTX is powered by Motr software, the software in CORTX that governs both object storage and key-value stores.

Motr

At the core of CORTX lies Motr. Motr is a distributed object-storage system that targets mass-capacity storage configurations. In order to ensure the most efficient storage utilization, Motr interacts directly with block devices rather than a local file system. The Motr design was heavily influenced by the Lustre file system, Network File System (NFS) v4, and database technology.

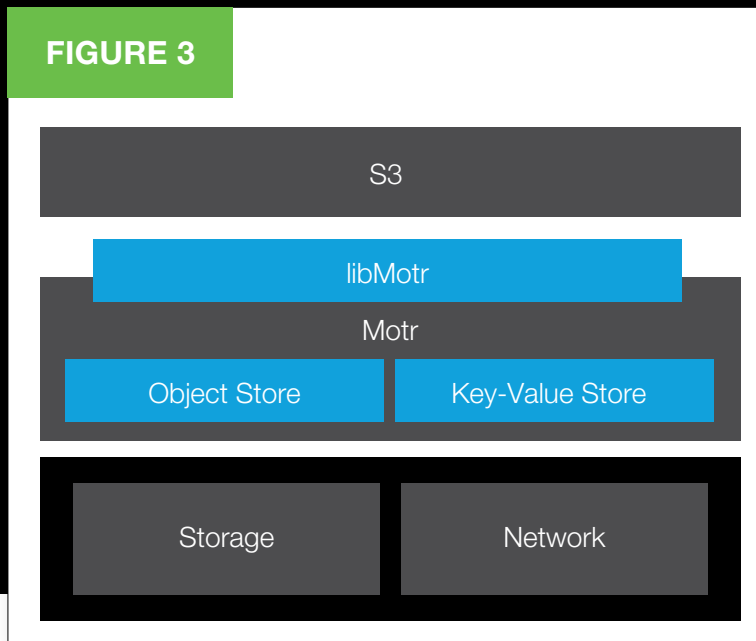
Traditional file-system properties such as hierarchical directory namespace and strong POSIX consistency

guarantees are no longer desirable or achievable at mass capacity. Instead, Motr is a more general storage system that provides an optional file-system interface. This approach allows a wider range of deployments, such as for the cloud.

Motr software provides CORTX with both horizontal and vertical scalability. Horizontally, Motr enables the addition of more CORTX nodes to increase storage capacity. It also enables handling extensions running on other nodes and helps avoid metadata hotspots by using a shared-nothing input/output (I/O) path. Vertically, Motr makes it possible to add memory and processors to servers running CORTX.



Figure 3. Simplified architectural diagram of the relationship between the S3 protocol and the storage network



Beyond scalability, Motr also provides fault tolerance, observability, extensibility, and portability. The flexible erasure coding in Motr takes the specific storage hardware and network topology into account, which helps enable fast network RAID repairs. Built-in monitoring in Motr collects information about system behavior for administrators, which provides observability. Motr also facilitates flexible transactions and an extension interface to increase extensibility. Finally, because Motr runs in the user space of any version of Linux, it also provides portability for CORTX.

The industry-standard S3 protocol allows multiple applications and frameworks, such as Apache Spark, PyTorch, TensorFlow, Apache Hadoop, Arrow, NoSQL, or Splunk, to use CORTX. The API in Motr provides a library of functions that applications and front-end programs can use for

accessing and manipulating storage resources using CORTX. The Motr API is an alternative interface to storage resources powered by CORTX, which provides fine-grained control over storage access.

Seagate Lyve Rack R1

Seagate Lyve Rack is a system-level integrated infrastructure solution available through Seagate partners, who test specific components together to deliver an enterprise-class object-storage solution. When organizations buy Lyve Rack through Seagate partners, they get enterprise-class object storage with CORTX pre-installed on fully tested, preconfigured hardware. They also get support, installation, deployment, and legal indemnification. Lyve Rack comes ready to deploy in the data center, without the high cost in time and money required to build an object-storage solution from scratch.

The base reference architecture for Seagate Lyve Rack is shown in Figure 4, which shows the management, data, and hardware layers of Lyve Rack. Because it is pretested and preconfigured, Lyve Rack helps solve many challenges around deploying object-storage software, and it simplifies the build-out of mass-capacity storage solutions. Lyve Rack is built using Seagate® Exos™ X 5U84 storage-area network (SAN) arrays that provide up to 1.3 PB of data in a single 5U enclosure,⁷ in addition to up to 99.999 percent high availability.⁸ Beyond the SAN arrays, the Supermicro servers—powered by Intel Xeon Silver 4210R processors—in Lyve Rack power the multi-tier data encoding in CORTX. This encoding provides better data protection, with rebuild times for failed disks that are up to 95 percent lower than with RAID.⁹ The Intel Xeon Silver 4210R processors that power Lyve Rack pair high performance with an excellent total cost of ownership (TCO).



Management Network

Lyve Rack provides single-pane-of-glass management across all nodes. As shown in Figure 4, all Lyve Rack nodes are connected to the management network, making for a redundant, highly available management layer. Seagate recommends a 1 gigabit per second (Gbps) or 10 Gbps network, as it is used for configuration and provisioning operations. Within the network, a single management virtual IP is exposed to customers; it is served by both nodes in the management network.

Data Network

Figure 4 shows how CORTX benefits from tighter integration with the storage. CORTX enables data to communicate directly with the storage drives, without an intervening file system. For the best performance, Lyve Rack R1 should be connected to a 50 Gbps network. As with the management network, a single data virtual IP is exposed to customers and is served by both nodes. Lyve Rack supports the S3 protocol.

Private Network

Server nodes are directly connected via a private network to support high availability (HA) and failover mechanisms. This network is also used for Motr internal communication between the nodes.

Storage Connectivity

The direct-to-drive architecture of Lyve Rack helps improve performance and reliability while granting better control over storage drives. Lyve Rack nodes redundantly connect to the storage hardware via Serial Attached SCSI (SAS) connectors for performance and availability.

FIGURE 4

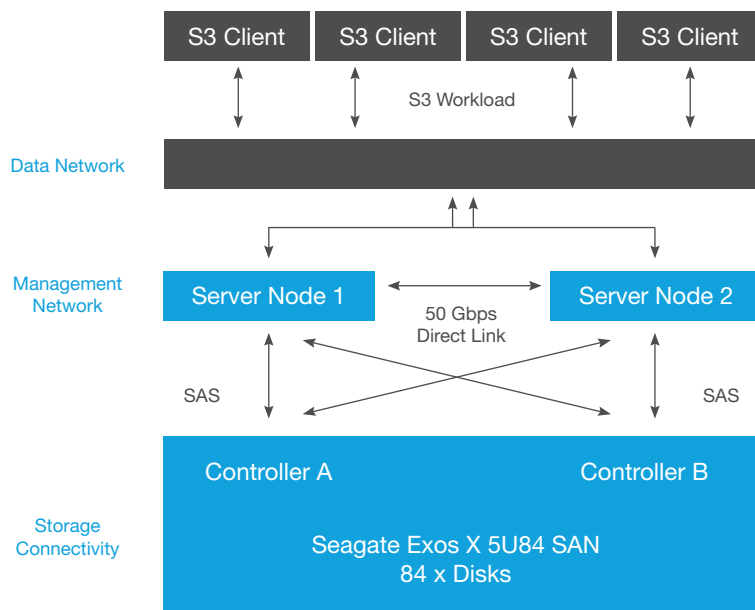


Figure 4. Simplified architectural diagram of Seagate Lyve Rack



Lyve Rack is optimized to use new disk storage innovations such as Heat-Assisted Magnetic Recording (HAMR)-based disk drives. Technologies like this increase hard drive capacity, increasing the density of bits stored on each square inch of surface space, further reducing the cost of object storage.

Affordably Deploy Object Storage

Seagate Lyve Rack delivers a simple, easy-to-deploy, high-capacity object-storage solution that you can use to manage data growth. Its validated, preconfigured reference architecture lowers both the cost and risk of adopting Lyve Rack and provides an efficient path to deploying CORTX open source software.

You can choose different partner products for your infrastructure while also having a choice about the native cloud protocols you use to easily manage and move data across storage solutions. Lyve Rack provides a building-block solution architecture for customers in need of mass-capacity storage. High-performance 2nd Generation Intel Xeon Scalable processors in Lyve Rack servers enable fast write, read, and search capabilities in Lyve Rack.

For More Information:

Contact Seagate sales: seagate.com/products/storage/object-storage-software/.

Visit seagate.com/products/storage/object-storage-solutions/lyve-drive-rack/.

Download CORTX and follow the growing CORTX community at <https://github.com/seagate/cortx>.

Visit intel.com/storage.

seagate.com

¹ 71 percent higher core count: Intel Xeon Gold 5220R processor core count = 24; Intel Xeon Gold 5120 processor core count = 14.

² 25 percent higher maximum Intel Turbo Boost Technology frequency: Intel Xeon Gold 5220R processor turbo boost frequency = 4.0 GHz; Intel Xeon Gold 5120 processor turbo boost frequency = 3.2 GHz. IDC. Based on using Seagate Exos 16 TB drives.

³ 11 percent higher memory speeds: Intel Xeon Gold 5220R processor DDR4 memory speed = 2,667 MHz; Intel Xeon Gold 5120 processor memory speed = 2,400 MHz.

⁴ 30 percent larger memory capacity: Intel Xeon Gold 5220R processor max memory capacity = 1 TB; Intel Xeon Gold 5120 processor max memory capacity = 768 GB.

⁵ 86 percent more cache: Intel Xeon Gold 5220R processor cache = 35.75 MB; Intel Xeon Gold 5120 processor cache = 19.25 MB.

⁶ IDC "The Digitization of the World." seagate.com/our-story/data-age-2025/.

⁷ Based on using Seagate Exos 16 TB drives.

⁸ Seagate. "Exos X 5U84 data sheet." May 2018. seagate.com/files/www-content/datasheets/pdfs/exos-x-5u84DS1982-1-1805GB-en_AU.pdf.

⁹ Seagate. "ADAPT: Rapid-Rebuild Technology." August 2018. seagate.com/www-content/solutions/backup/backup-and-recovery/SB505-1-1808US_Adapt-Technology.pdf.

© 2021 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. AcousticShield, Exos, the Exos logo, Seagate Secure, and the Seagate Secure 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. 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. Actual data rates may vary depending on operating environment and other factors, such as chosen interface and drive capacity. Seagate reserves the right to change, without notice, product offerings or specifications.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.



SEAGATE