



AVAILABILITY AND DATA DURABILITY

Company Overview

Seagate® crafts the datasphere, helping to maximize humanity's potential by innovating world-class, precision-engineered data management solutions with a focus on sustainable partnerships. Seagate has been a global data management leader with more than 40 years of innovations. Lyve™ Cloud is the next phase of evolution for Seagate, delivering storage as a service that complements existing cloud offerings and fills the unmet needs of mass-data storage.

Executive Overview

Lyve Cloud is a simple, trusted, and efficient storage as a service for mass-capacity storage. Predictable economics with verifiable trust and ease of use at scale make Lyve Cloud the right choice to put mass data to work. More data can be stored for longer and activated whenever for whatever application.

With Lyve Cloud S3 storage as a service, enterprises can protect and store any amount of data for a variety of use cases. These include data lakes, backup and restoration, archives, enterprise applications, IoT devices, and big data analytics in an object storage platform.

Lyve Cloud provides a highly efficient, durable, and reliable object storage that is crafted to preserve data integrity and ensure the service is resilient and fault tolerant. As with any cloud service, a key aspect to consider in addition to the core product is data accessibility and usability. This document is focused on the availability and durability aspects of the service and how we achieve this for our customers. It also focuses on the benefits of having data closer to where it is generated and used.

Seagate is confident in its Lyve Cloud S3 storage as a service's ability to meet the high-level service level agreements (SLAs) of its customers regarding data availability and durability. This is further supported by Lyve Cloud's transparent, simple pricing structure and predictable economics. With no ingress or egress fees and zero API charges, customers have the freedom to move their data wherever they want if they aren't satisfied with Lyve Cloud's availability and durability.



Simple



Trusted



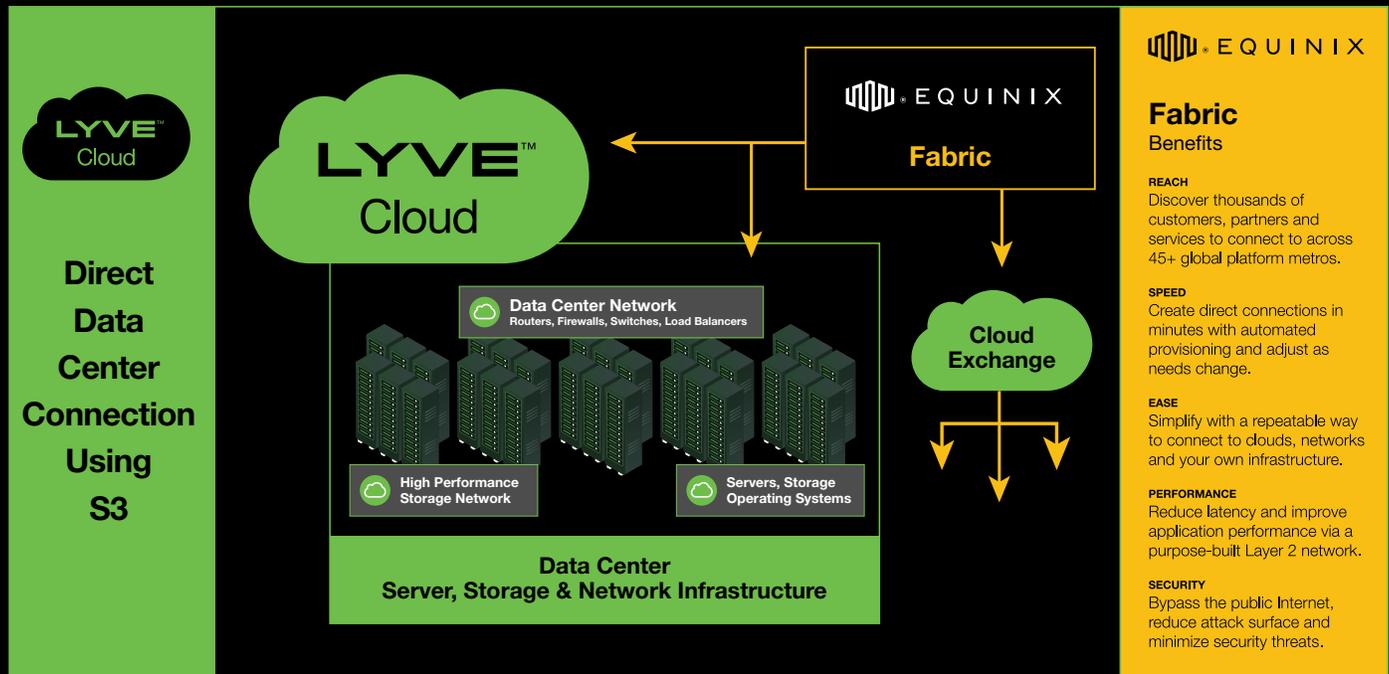
Efficient



Data Accessibility & Availability

Lyve Cloud services are offered at Equinix data centers, which are highly secure, fully redundant, and certified for SOC2 and ISO27001 compliance. Design of the infrastructure elements includes 1+1 redundancy.

Lyve Cloud doesn't stop at data durability. Our collaboration with Equinix assures data is also highly available. By locating with Equinix, Lyve Cloud gives customers a direct connect via ECX Fabric or Metro Connect. This supports high-speed access to Lyve Cloud without the need for increased internet bandwidth.



The presence of Lyve Cloud at Equinix metro edge locations also brings customers closer to where data is being generated and used. It also brings the possibility of direct connectivity to all major cloud service providers.

Lyve Cloud offers geographic replication capabilities that allow the customer to replicate data to other regions, fulfilling their compliance and regulatory requirements or ensuring continuous service in the event of a regional catastrophic failure.



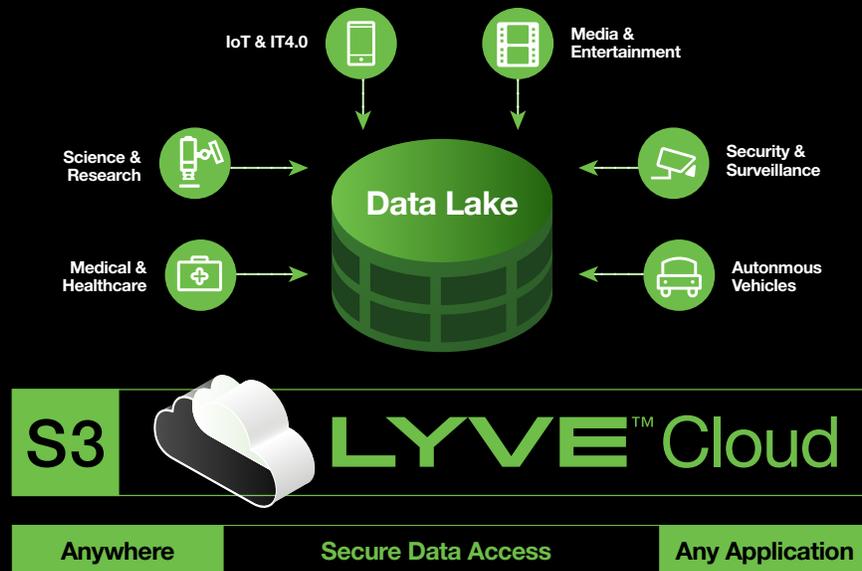
Object Storage & Erasure Code

From multimedia files and text documents to web pages and log files, traditional hierarchical file storage systems and block storage are not always the best fit for these large volumes of unstructured data. Object storage is the alternative to file- and block-based storage for big data, as demonstrated by organizations with hyperscale environments. S3 makes object storage simple to consume and integrate, as it's delivered as a service over the web.

S3 object storage provides the right combination of cost-effective scalability, data integrity, and accessibility required by the modern enterprise. Each object has metadata and uses the URL as a unique identifier. This eliminates the need to know the exact location of data within the storage environment. Every object is accessible from anywhere through its unique URL, which enables maximum flexibility for data access or distribution from anywhere.

S3 object storage offers the flexibility to keep storage separate from compute resources, allowing businesses to optimize their compute and storage, as well as network and applications. Users can consume their S3 storage resources in the optimal way for each particular workload.

With this benefit, S3 object storage is quickly becoming an integral element of enterprise storage infrastructure and digital transformations, as it's designed to scale for capacity and performance. Unlike hierarchical file systems with files and folders, object storage uses a flat structure. This enables users to store billions of files without the complexity and performance issues that can develop in hierarchical file-system environments.

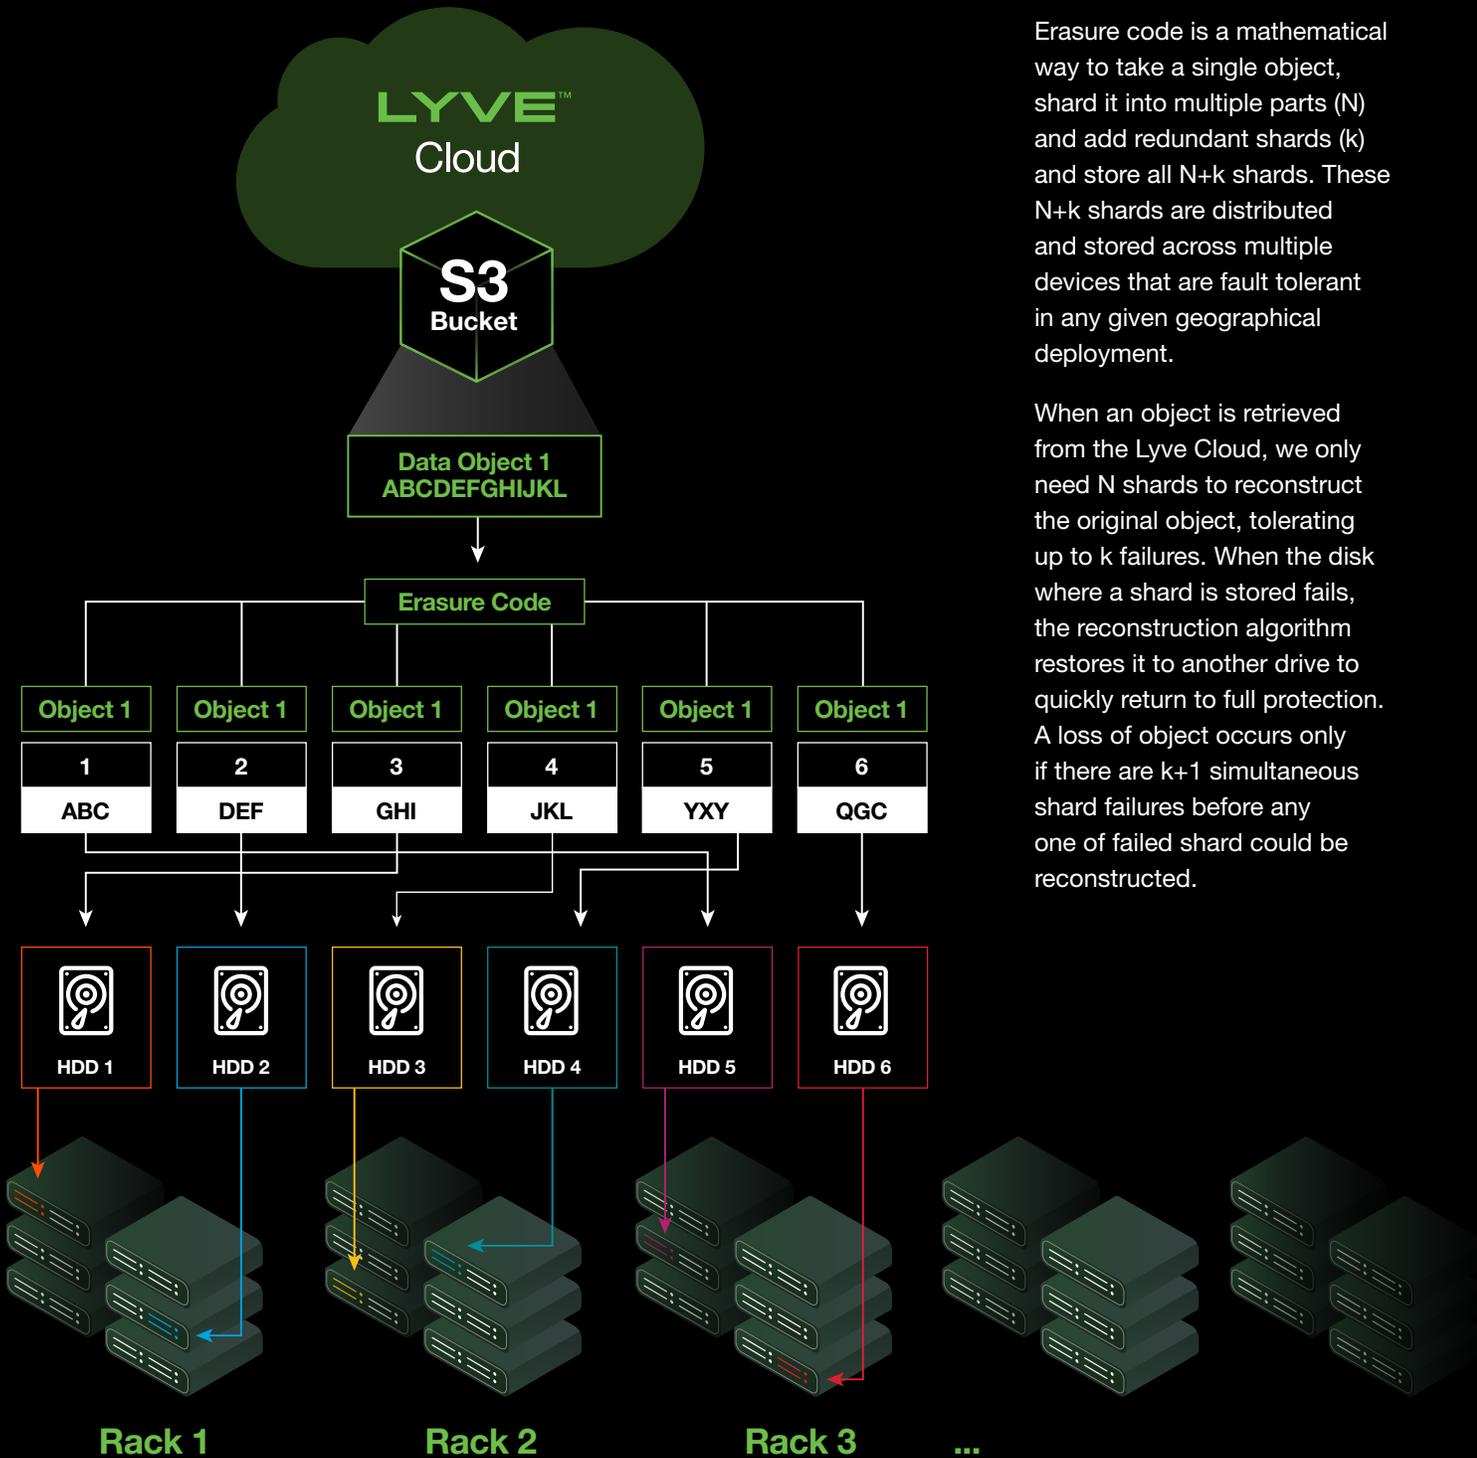


Erasure Code and Data Durability

Lyve Cloud is an object storage system, and its storage nodes are based on erasure code for high data availability, durability, and performance in scale-out storage applications.

Data Durability is how an object storage platform is gauged. Lyve Cloud provides eleven 9's (99.999999999%) Data Durability. This is achieved by leveraging industry standard Reed-Solomon Algorithm based erasure code.





Erasure code is a mathematical way to take a single object, shard it into multiple parts (N) and add redundant shards (k) and store all N+k shards. These N+k shards are distributed and stored across multiple devices that are fault tolerant in any given geographical deployment.

When an object is retrieved from the Lyve Cloud, we only need N shards to reconstruct the original object, tolerating up to k failures. When the disk where a shard is stored fails, the reconstruction algorithm restores it to another drive to quickly return to full protection. A loss of object occurs only if there are k+1 simultaneous shard failures before any one of failed shard could be reconstructed.

This is a conceptual drawing to explain how object storage is using erasure code and distributes object shards on storage nodes distributed in different racks for Lyve Cloud. It does not reflect the actual setup and erasure code methodology.



Infrastructure and Storage Innovation

Lyve Cloud is built upon the infrastructure foundational principles of being highly available, fault tolerant, and resilient. Hosted at Tier 4 data centers at the metro edge, Lyve Cloud enables high performance with low latency. As a leader in the data management space, we continuously innovate to push the boundaries on hard drive (HDD) capacity, reliability, and efficiency. Taking advantage of these innovations, we deliver a robust cloud storage platform that provides a reliable and resilient service.

Infrastructure and Storage Resiliency

It's important to understand the factors that contribute to Lyve Cloud's resiliency. Both the number of HDDs that data is stored in and the number of storage nodes these HDDs are spread across matters—as well as where these nodes are located. Having all the storage nodes within two or three data center racks mean that a single rack failure could impact access to the entire storage on those nodes, resulting in data inaccessibility—or worse, data loss.

In line with our foundational principles, the data shards are not only spread across multiple number of storage nodes, but each of these nodes are physically in separate data center racks. As an example, our units comprise 16 individual storage nodes, each in their own rack. Data is sharded across all the storage nodes, and any individual failures within the infrastructure—be it a disk, storage node, cabinet, switch, or several of these failures—will not impact data accessibility.

Fault Tolerance

Each of the storage nodes are comprised of several HDDs. Lyve Cloud's erasure coding algorithm places the shards in HDDs of the storage nodes, ensuring no two of the same object is on the same HDD or storage node. As a result, it would take more than 50% of our complete infrastructure to fail before any data could become inaccessible. To put this in perspective, that would be the equivalent of simultaneous failure of thousands of HDDs or dozens of servers.

The chance of multiple simultaneous HDD failures in any given deployment is not uncommon. If the probability of having a non-recoverable sector in a hard drive happens at 1000 terabytes (TiB) of data written and read from a disk, having a second or third failure while data is available and accessible is a real possibility. We mitigate this risk using predictive failure algorithms that keep track of sectors that could potentially fail. In this case, data is progressively relocated to another disk drive.

Data Integrity

Silent data corruption (SDC), otherwise known as bitrot, is often not immediately visible and is only detected when corrupted data is requested. Periodical data integrity scans are a common practice, but this passive approach doesn't eliminate the problem. We actively track and fix SDC/bitrot by calculating a checksum for every sector on the HDD and storing it for verification.

Lyve Cloud also leverages industry standard SAS channels and protocol to connect storage and disk systems. A sustained data rate of 1TiB/s continuously for one year yields 0.3 SDC events per SAS channel, a non-recoverable error rate of 10^{21} , which is significantly higher than SATA low-cost storage systems. This is further improved upon by Seagate's innovation to the SAS controller—adding protection information according to T10 standards and enhancing SDC to 10^{28} .

Data Durability

Lyve Cloud offers 11 nines of durability. Since there isn't a standard industry model to calculate this, and to avoid any discrepancies, we've applied the same approach as other providers to compare our durability numbers. This estimate is based on disk drive failure rates and restoration times. Annual failure rate (AFR) is usually 0.35%, and for the sake of keeping in line with the comparison to another provider, we've used a value of 5%.

Disk Drive repair times are estimated by multiplying usable disk capacity by write speeds dedicated for rebuild, which comes down to a mean time to repair (MTTR). With Lyve Cloud, however, there isn't such a thing as MTTR that would impact data durability. That's because Lyve Cloud uses erasure code, which would rebuild and relocated data from a defective disk drive



