



A NEW APPROACH TO CLOUD STORAGE

Lyve™ Cloud is a simple, trusted, and efficient on-demand solution 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 any application.



CONTENTS

- 3** THE LANGUAGE OF INTERNET DATA STORAGE
- 3** WHAT ARE S3 BUCKETS FOR?
- 4** S3 OBJECT STORAGE VS. TRADITIONAL STORAGE
- 5** USE-CASES AT THE METRO EDGE
- 6** A COMPLEMENTARY CLOUD STRATEGY
- 7** BEST-FIT APPLICATIONS
- 7** BACKUP AND RECOVERY
- 9** ACTIVE ARCHIVES
- 10** STREAM LOG FILE REPOSITORY
- 11** VIDEO SURVEILLANCE




The Language of Internet Data Storage

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 is a cloud-based object-storage-as-a-service solution that is fully S3 compatible. The Amazon S3 application programming interface (API), otherwise known as Amazon's Simple Storage Service, can be thought of as the language of internet data storage. Lyve Cloud combines object storage with standard S3 language. This makes it easy to connect internet-based applications for data storage and activation. It also enables Lyve Cloud to be used alongside private, hybrid and multicloud data centers.

With S3, data is stored in objects that are grouped into logical containers called buckets. Each object is identified by a unique, user-assigned key. Buckets can be managed by the client administrator using either the console interface or via the API that's integrated in applications that communicate with and use S3 storage services. Users can simply create a bucket in the Lyve Cloud management portal and submit the unique bucket name and credentials that were created in the application software.



What are S3 buckets for?

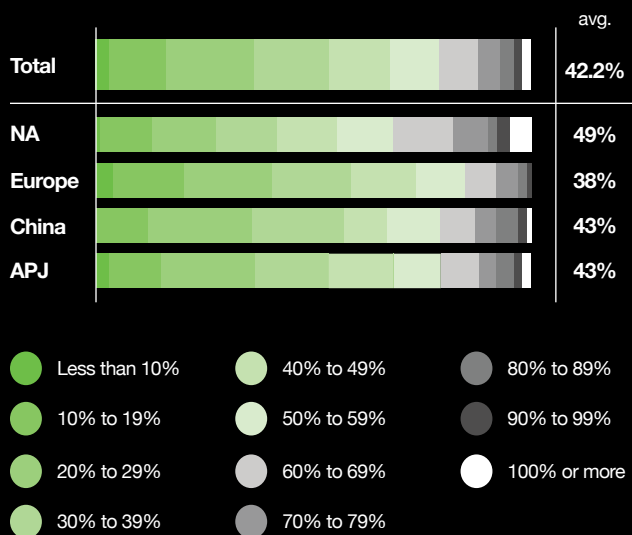
An S3 bucket is a data repository that's similar to a file folder, but unlike a file folder there is no hierarchy and no limits. S3 buckets store objects, which consist of data and its descriptive metadata. These buckets help users organize data stored in an S3 cloud and allow them to specify different rules and policies within their data management workflows.

Object storage delivered as a service with an S3 interface is the optimal choice for accommodating fast-growing volumes of unstructured data. Organizations can also achieve the robust performance required for supporting computationally intensive workloads, including artificial intelligence (AI)/machine learning (ML), Internet of Things (IoT), and big data analytics.

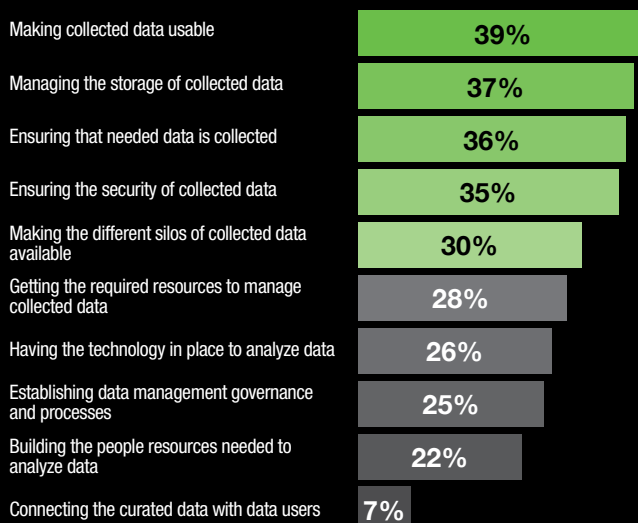


Organizations across industries are launching new technology initiatives that require them to store, access, and analyze large, fast-growing volumes of data. Much of that data is unstructured, whether organizations are implementing AI/ML, IoT, or other big data solutions. As such, these users need an effective way to store and analyze terabytes, petabytes, or several petabytes of data.

Expected Annual Data Growth Rate



Barriers to Putting Data to Work



Source: The Seagate Rethink Data Survey, IDC, 2020

S3 Object Storage vs. Traditional Storage

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.

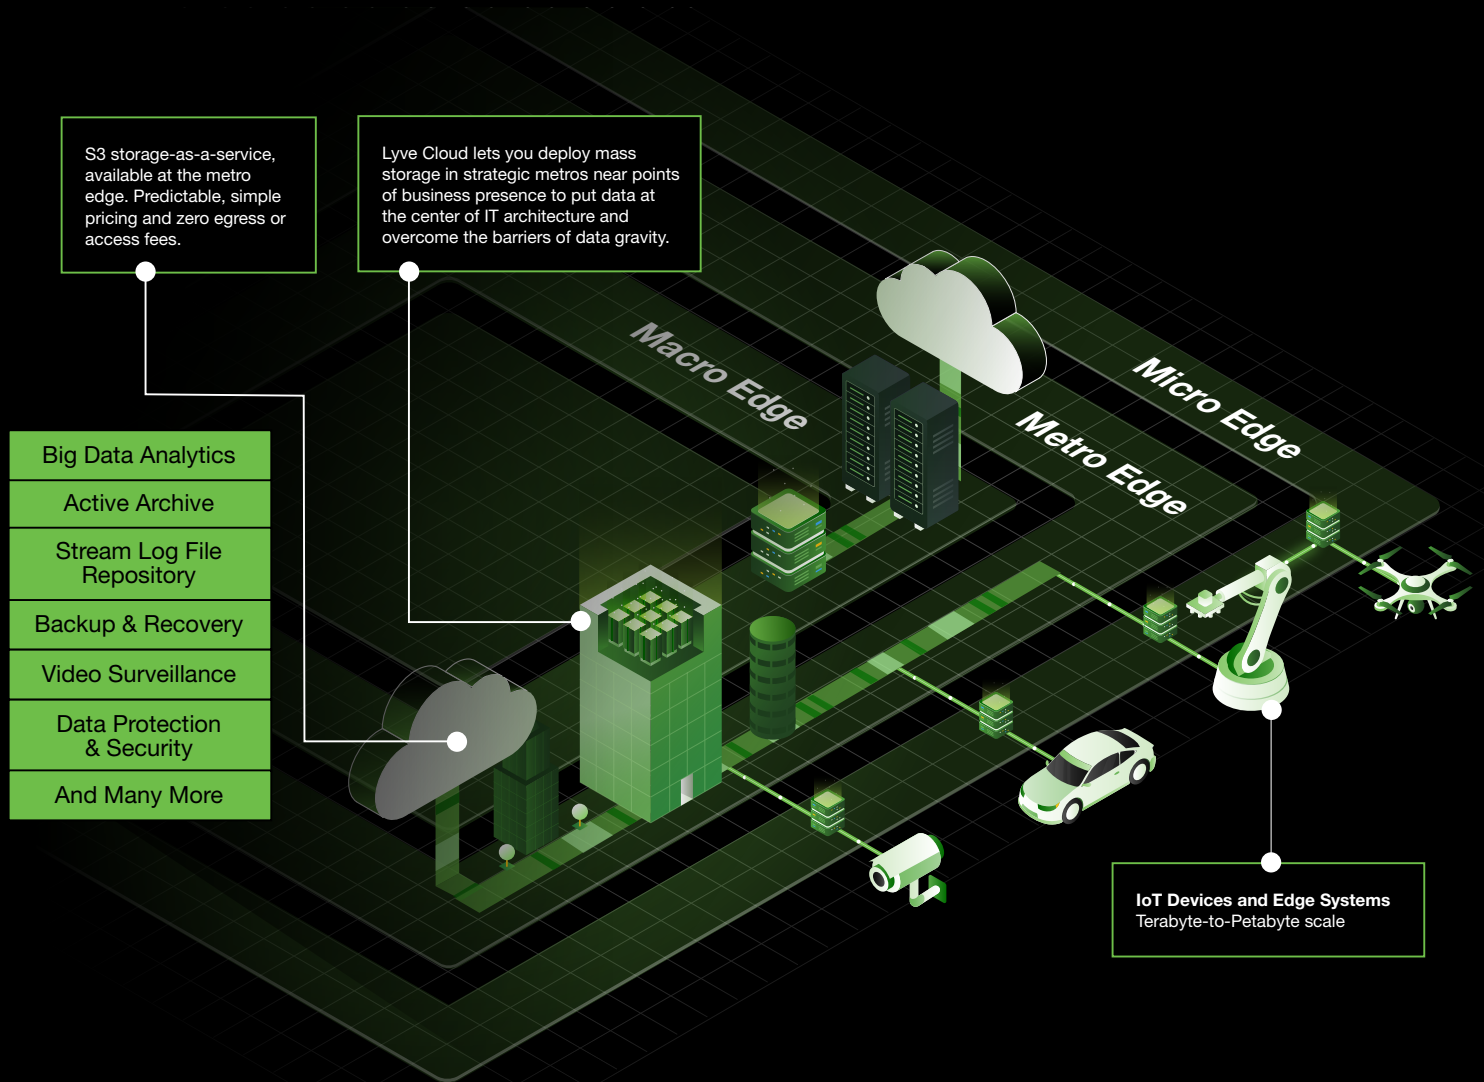
S3 object storage offers the flexibility to keep storage separate from compute resources, allowing businesses to optimize their compute and storage, as well as network and applications. Users can consume their S3 storage resources in the optimal way for each particular workload.

With this benefit, S3 object storage is quickly becoming an integral element of enterprise storage infrastructure and digital transformation, 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.



Use-Cases at the Metro Edge

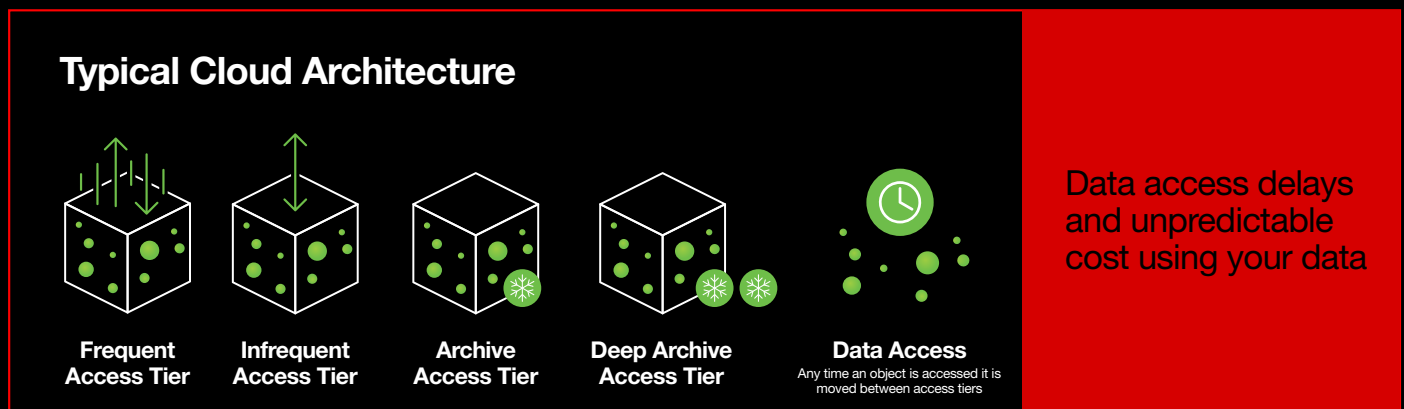
Take control of your data with Lyve Cloud. Interconnect with any cloud or compute ecosystem to enhance multicloud performance and efficiency.



A Complementary Cloud Strategy

Lyve Cloud S3 storage complements existing hybrid and multicloud environments with the standard S3 API integration. S3 storage can be employed to store any type of object, which enables uses such as storage for internet applications, backup and recovery, disaster recovery, data archives, and data lakes for analytics.

High-performance S3 storage allows organizations to unlock the full value of their data with scalability, accessibility, data protection, and the cost-effective nature of object storage. Unlike other public S3 cloud storage services, Lyve Cloud delivers always-on data in a self-contained storage tier for best performance and optimal total cost of ownership (TCO). Below is an example of the storage tiers that are offered by typical cloud providers.



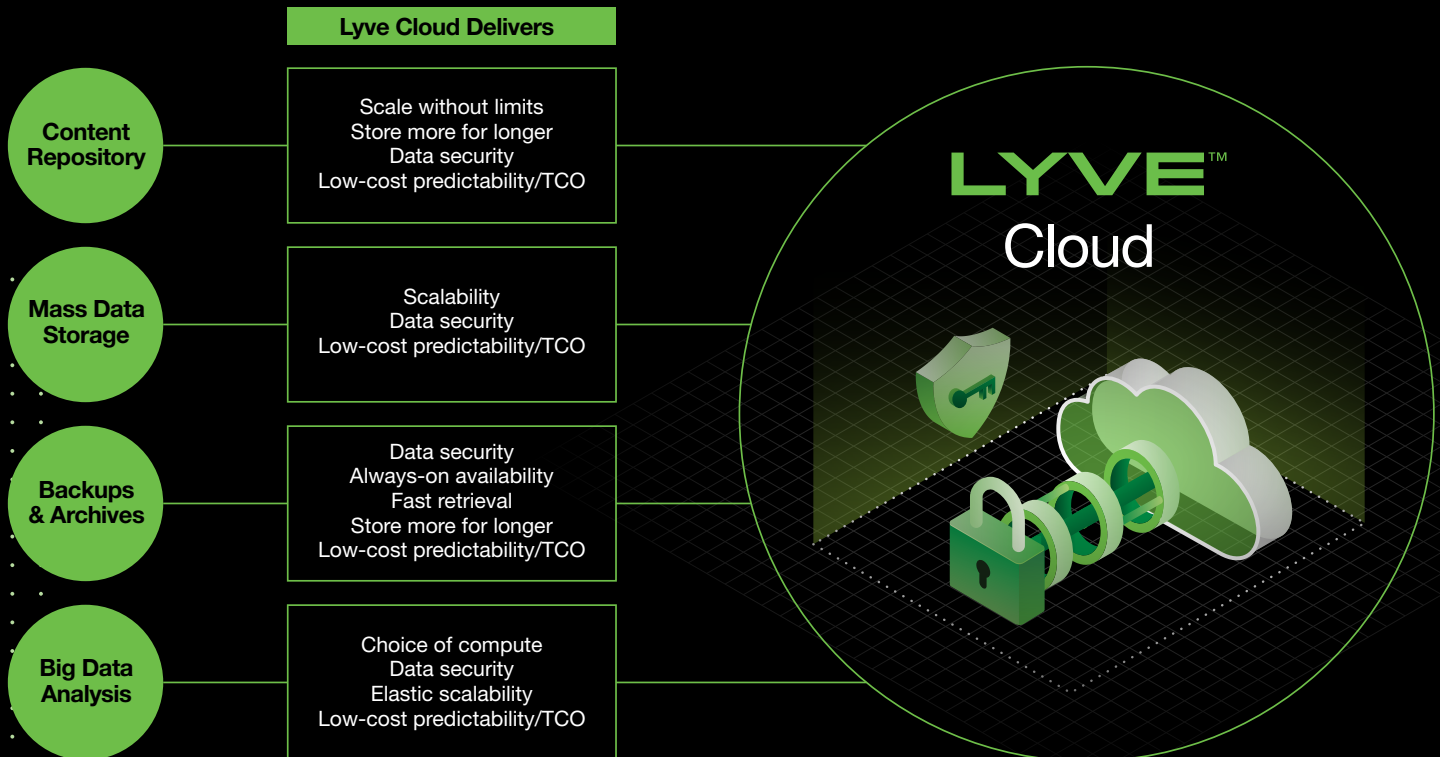
With Lyve Cloud, data is available when it's needed without having to wait or move to higher, faster, and more responsive storage tiers. Exabytes of data can be stored and activated with predictable economics. That's because pricing is based on capacity stored, and there are no additional charges for API calls or egress fees to use data or move data in and out of S3 buckets.

Simple scalability and long-term cost predictability make Lyve Cloud a great complement to existing cloud strategies for a variety of use cases. Enterprises can store mass data captured in daily operations and activate it for their business.



Best-Fit Applications

Lyve Cloud provides secure long-term retention, fast retrieval, and industry-standard data protection. Users benefit from a flexible and highly scalable platform that allows them to store more of their data while working seamlessly with their choice of compute and applications. Lyve Cloud is best suited for the four storage use cases highlighted below.



Backup and Recovery

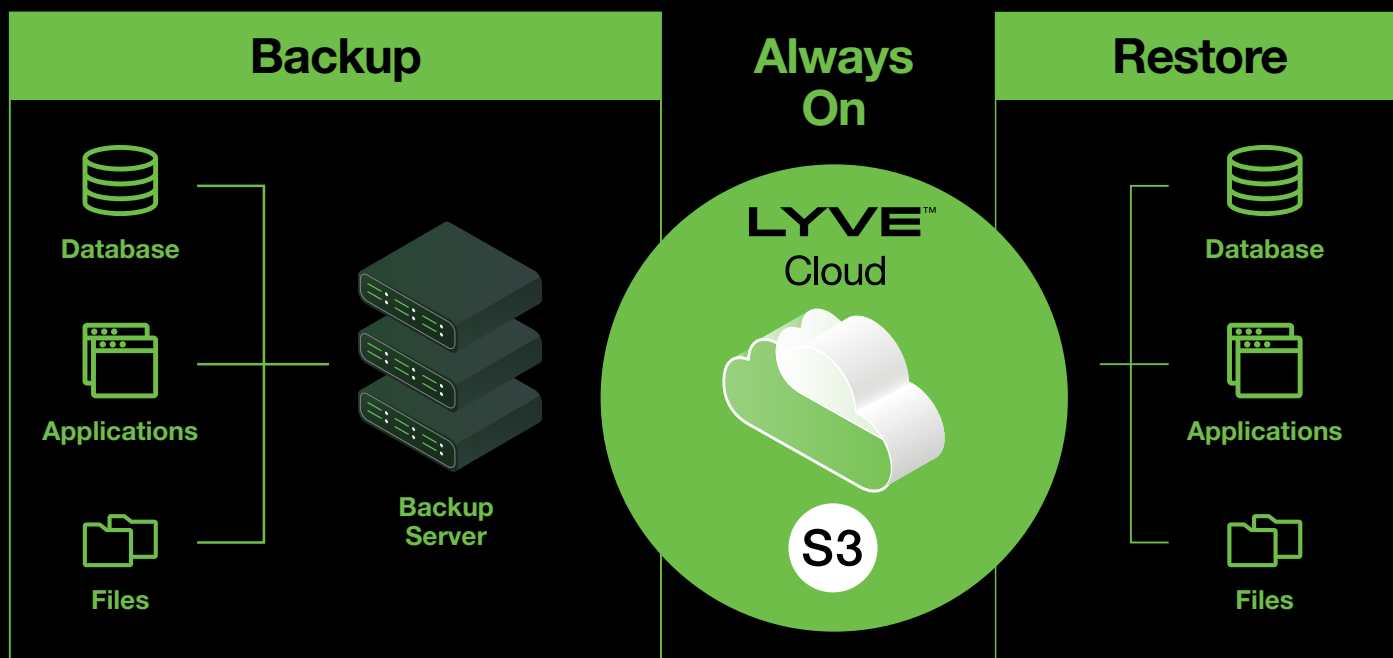
Lyve Cloud provides durable, secure, and cost-effective cloud storage for a variety of backup needs. With Lyve Cloud, data is always on for backup and recovery. Restores can't wait when it comes to business continuity!

Data protection for business continuity is no trivial task. At a high level, and according to best practices, businesses should follow the 3-2-1 rule. This rule calls for having at least three copies of data on two separate storage devices, as well as one copy stored offsite at a secure location. Enterprise backup environments have grown organically over years. Now, these environments tend to have multiple different backends for storage with backup to disc to tape. Depending on an enterprise's recovery time objective (RTO), Lyve Cloud can streamline backup operations while simplifying the archive process.



To ensure business continuity, organizations can replace their tape backups and archives by using Lyve Cloud as a co-location data center for disaster recovery. This backup data can also be reused for secondary use cases such as analytics and activation, leading to better insights and operational improvements. By setting backup and archive data as immutable, businesses can prevent their data from being manipulated. This will protect the data from ransomware and accidental deletion while assuring availability in the event of a compliance audit.

Most backup software will integrate seamlessly with Lyve Cloud using the S3 protocol. This allows enterprises to centralize data in a single scale-out backup repository, which puts less of a burden on existing storage infrastructure for backup while reducing cost and complexity. Long-term mass-data storage at a compelling TCO means worry-free business continuity.



Businesses can use Lyve Cloud to:

- Replace costly and inefficient tape backups without worrying about tape-to-tape data migration, on-premises tape-library maintenance, or offsite physical archiving tasks. S3 buckets and object immutability safeguard data against ransomware attacks while achieving the same air-gap protection as tape when used for disaster recovery.
- Scale or replace primary on-premises backup repository with high-performance S3 storage that can easily meet RTOs for daily restores and disaster recovery. Businesses pay only for the storage capacity they need, which lowers TCO.
- Keep their backups separate from applications and data by storing them securely in a different data center location.
- Backup their application databases and fast-growing file server.
- Activate their backup data with analytics tools.




Active Archives

Old yet potentially useful data is stuck in legacy environments or scattered across different locations in data silos. Lyve Cloud enables businesses to move and ingest mass data into a single cloud, which makes it easier for them to access their data and use it to glean new insights that improve revenue and operations.

S3 storage transforms traditional archives into active archives, thereby assuring enterprises their data is active, available, and preserved. By reducing their IT infrastructure, businesses will also reduce operating costs and the complexity of their storage. This is in addition to making data that was previously out of reach more accessible, usable, and manageable for the business.

Enterprises can leverage Lyve Cloud for their digital transformation, using it to build a data lake so they can future-proof operations across different business disciplines. Lyve Cloud enables enterprises to store as much mass data as they need and to use it whenever they need it—all without access fees or costly delays. Further, it supports real-time data processing when combined with analytics while preserving data for future use and compliance.



Lyve Cloud supports:

- Flexible data archiving that makes data accessible from anywhere, anytime, and for any application using S3.
- Converting traditional tape archives to active archives and eliminating maintenance costs. Archiving to S3 storage with Lyve cloud doesn't incur upstream costs.
- Mass-data retention for data-intensive environments such as healthcare, IoT, media and entertainment, and surveillance.

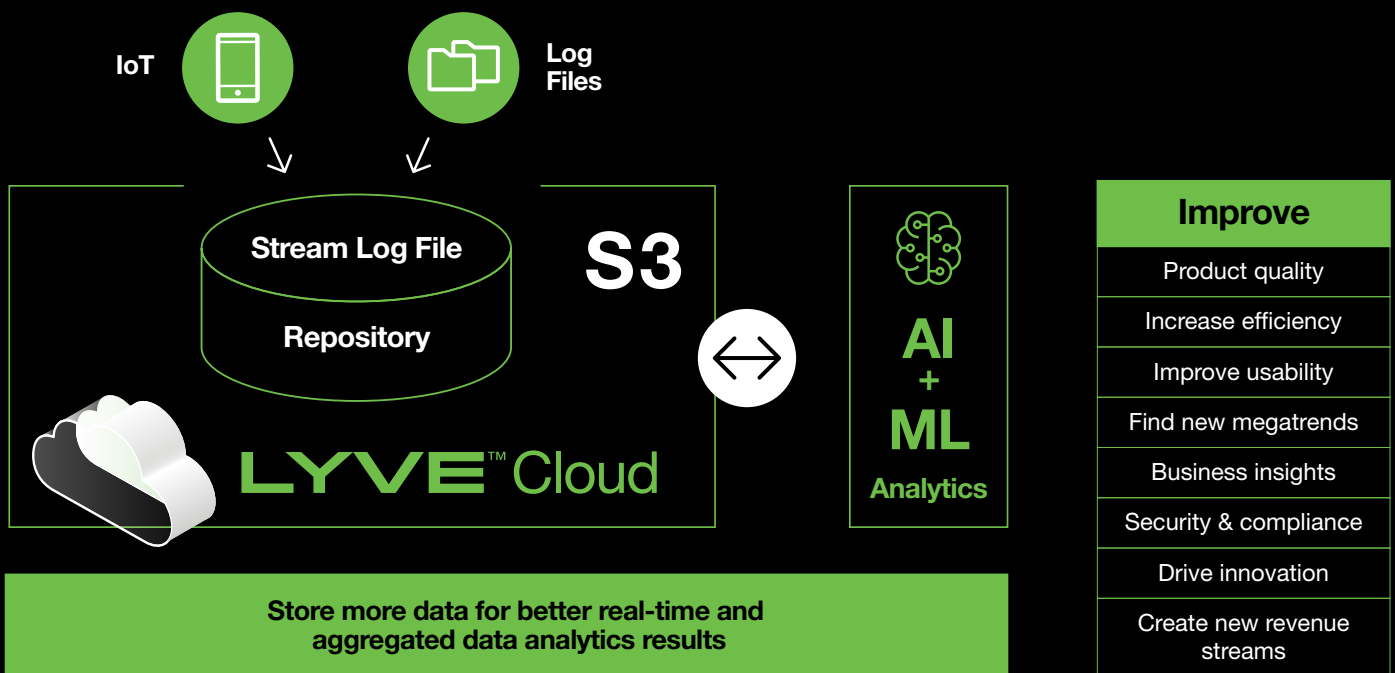


Stream Log File Repository

Data growth is exacerbating data gravity at the metro edge. One of the biggest contributors to this explosive data growth outside of raw data produced by applications is log-file data. This log-file data holds tremendous value, but storing and analyzing all of it remains a challenge. With so many consistent streams of log-file data being generated daily—particularly in IoT environments—businesses need dynamic cloud storage with predictable economics to maximize the data's full potential.

Some log files must be analyzed in real time while others must be aggregated in a pool for big data analytics. Collected logs must also be preserved for compliance and security audits. It's a very dynamic data environment with frequent updates and changes. As such, storing log-file data in a cloud storage bucket would cause significant and unpredictable costs. That's why log-file data analytics would benefit from scale-out cloud storage solutions that can store data more dynamically than it can be stored on premises in a data center.

Lyve Cloud helps businesses overcome log-file storage and retention barriers by eliminating the cost uncertainty associated with storing data in the cloud and applying analytics as needed. A flexible, vendor-agnostic solution, Lyve Cloud is crafted to complement existing cloud ecosystems while delivering optimal TCO and world-class data security.



Video Surveillance

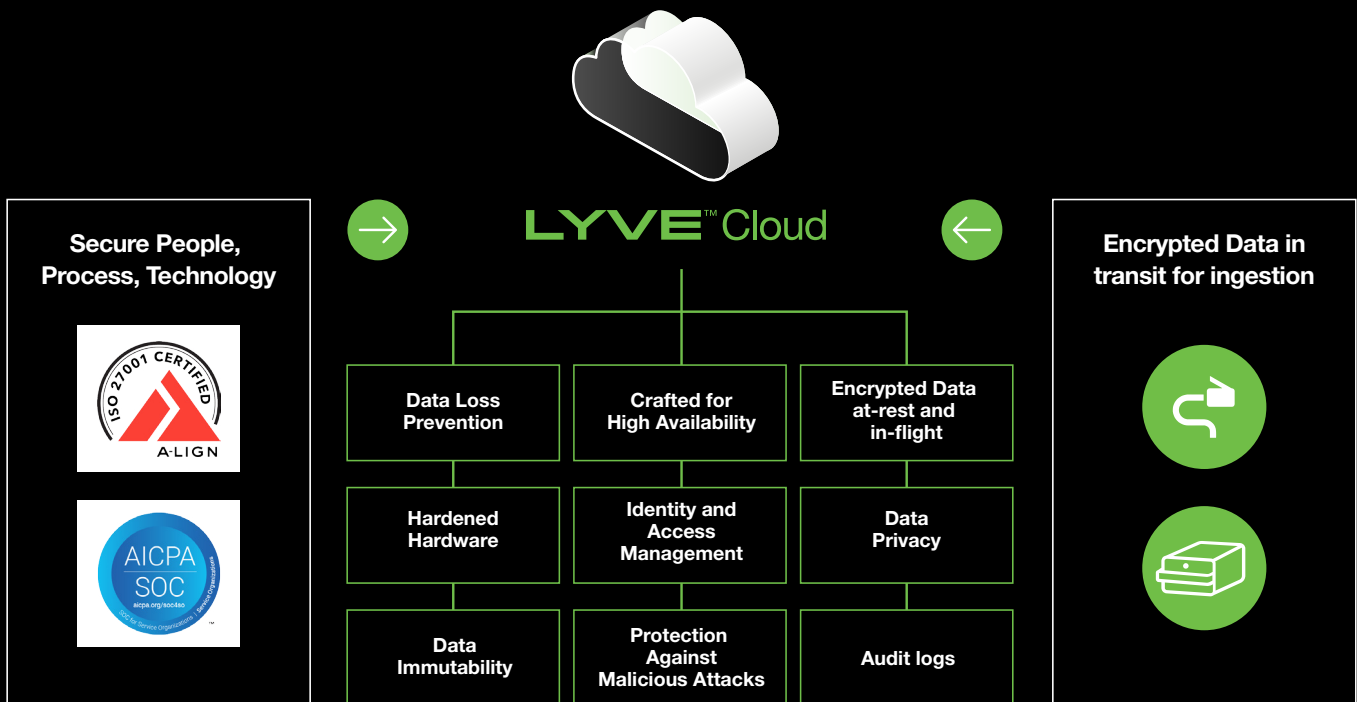
Lyve Cloud simplifies long-term video storage. How long surveillance video data must be stored depends on industry-specific regulatory compliance mandates. For some industries, such as banking, this means holding onto petabytes of video surveillance footage for half a year or more. When deployed as video surveillance as a service, or VSaaS, Lyve Cloud can help businesses overcome their video storage retention challenges.

Lyve Cloud gives businesses the flexibility to complement onsite video management systems or establish an entirely cloud-based solution. Businesses can choose to send video surveillance data collected onsite to Lyve Cloud after a specified duration of time. Video footage can be stored directly to and accessed from Lyve Cloud. This can all be accomplished with zero API extras or egress charges, which means businesses will pay only for the storage they use.

Enterprises can store more video surveillance data securely in Lyve Cloud, utilizing end-to-end data encryption for data in flight and at rest. Audit logs and S3 object Immutability prevent data from being modified, manipulated, or deleted, enabling users to specify retention times. Lyve Cloud's adherence to the most stringent global security standards is reflected in its ISO 27001 and SOC Type-2 certifications.

World-Class Security

Lyve Cloud adheres to the most stringent, globally recognized data security standards.



Conclusion

S3 is the universal interface and language used to connect object storage through the internet. When businesses use S3-compatible storage as a content repository, they can overcome challenges related to on-premises storage and workload management for a variety of applications. S3 enables the consolidation of mass data storage,

backup, and archive in a single, scalable object storage repository, making it easier to apply analytics tools.

Lyve Cloud is designed to complement your existing cloud strategy, offering true cloud flexibility. Further, it's a vendor-agnostic solution that helps you avoid lock-in while providing choice of compute

for S3 workloads. A simple pricing structure with long-term cost predictability means you'll never be surprised by your cloud bill. Put your data to work with always-on availability, world-class security, and cloud flexibility from the global leader in mass data storage management.



Ready to Learn More?

Visit us at seagate.com/services/cloud/storage
Or [download the brochure](#)

seagate.com

© 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. Lyve is either a trademark or registered trademark 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. Seagate reserves the right to change, without notice, product offerings or specifications. TP728.1-2104US April 2021

