SEAGATE
LYVE™ Cloud

EXPERIENCE MULTICLOUD THE WAY IT WAS MEANT TO BE

SEAGATE

CONTENTS

More Data Calls for More Cloud

Data volume is growing daily at organizations around the globe. IDC predicts that the collective sum of the world's data will grow to over 220ZB by 2026.[1] The Seagate® *Rethink Data* report shows that, just in the space of the next two years, enterprise data is expected to grow 42.2% annually. In lockstep with that growth is the need to harness data for analytics and intelligence that drive innovation and improve business outcomes.

Which brings us to how companies achieve these outcomes—the cloud. A single, ubiquitous term that encompasses a vast landscape of unique storage architectures with equally unique benefits and challenges. The cloud we'll focus on in this paper—and the one that a growing number of organizations are turning to—is multicloud.

Multicloud relates to deploying cloud environments from more than one cloud vendor with the same architecture. So for instance, you can have two private clouds from two vendors. Or, two public clouds from two vendors. This is contrast to deployments such as hybrid cloud, which use both public and private cloud vendors.

Seagate's *The Multicloud Maturity Report*, which provides insights from a global web survey administered in early 2022 of senior industry leaders, revealed that 82% of the 500 respondents already use 2+ public cloud infrastructure service providers (excluding Software as a Service, or SaaS)—and this figure is expected to increase to 93% in two years. Suffice it to say, multicloud is on the rise.

Here we'll examine the merits of multicloud, the challenges it poses for businesses, and how organizations can harness this ecosystem to unlock their data's full business value.
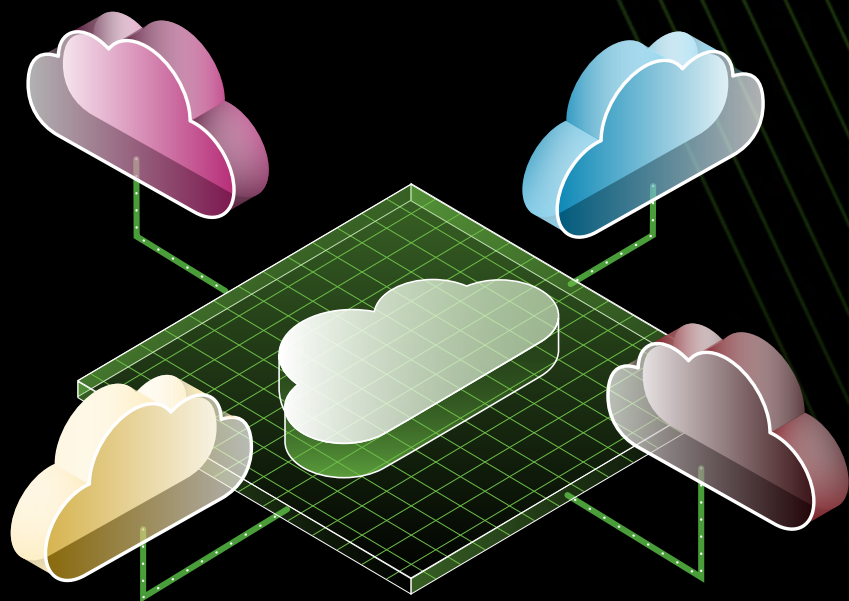
Multicloud Freedom

The Promise and Pitfalls of Multicloud

The multicloud is attractive for good reason. It allows organizations to match solutions to business and organizational needs by combining the scalability and flexibility of various public cloud services with the predictable economics of scale that private cloud can provide. When executed successfully, businesses can benefit from increased data protection, flexibility, and accessibility. The problem is successful execution is tough because of two primary constraints. Complexity and cost. Let's talk about each.

As data continues to boom and organizations do their best to wrangle it, much like a spider web, multicloud ecosystems are growing ever more expansive and intricate—and, therefore, ever more complicated. The Multicloud Maturity Report illustrates this reality by revealing that 56% of organizations storing hot/warm data on cloud infrastructure say that data is accessed by applications that run in a different environment on at least a daily basis. More than half of the respondents say they manage 100+ intercloud integrations. And not surprisingly, 76% say monitoring, measuring, and ensuring SLA adherence for applications that rely on these intercloud integrations is challenging.

All of this complexity equates to one thing: a multicloud fog. Where data is too often siloed, workload portability is limited, compliance risks abound, and companies are forced to double-down on security and resiliency to minimize weak links.

And it doesn't stop there. Another result of this complexity is the often inscrutable and volatile data-related costs, which further intensify friction. Most cloud repositories have raised up tall walls, making flexible data movement nearly impossible. When data wants to get in—that's easy. But leaving is another story. Between storage media costs, high egress fees, unpredictable API charges, and ambiguous pricing models, companies struggle to move data to where it's needed, or where it can render the most business value. In response, a vast majority of businesses end up deleting otherwise valuable data.

The consequence is that businesses using the multicloud often face a lose-lose-lose situation: They keep all their data but don't actually use it because of high costs. They make full use of their data but pay massive fees that slow or halt other areas of the business. Or they get rid of data (since they can't use it anyway) and risk losing out on potential insights and business value.

This is an unsavory set of options for any organization that wants to remain competitive. The good news is it doesn't have to be this way. There are three key storage issues that, when addressed and optimized, can free your data from the confines of a multicloud mess.

First: Control Costs

Data is critical to business strategies. IT and business leaders need to extract value from legacy data and seamlessly create, store, move, and analyze new data to reveal insights and create new revenue streams. But TCO is a challenging issue when it comes to storing data in the cloud. Costs to consider include overall storage, read access, write access, list and create operations, data retrieval, and support. In the multicloud, where exorbitant API charges and egress fees dominate, IT leaders are faced with trying to budget and innovate while dealing with wildly unpredictable bills and consistent overages.

At best, this kind of inefficiency stifles growth. At worst, it kills it. Therefore it's critical that organizations, especially those that already have or are scaling towards petabytes of data, search for cost-effective solutions that exclude API charges and egress fees, and that offer clear, consistent, predictable pricing.
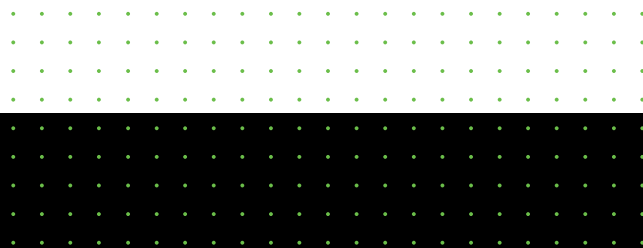
Second: Find Flexibility

Companies that innovate the most, and the fastest, in the multicloud have three things in common:

- They can access their data anytime, using and sharing it as needed.
- They're able to easily move data to wherever it'll be most valuable.
- They can quickly scale, storing more or less data as business needs evolve.

In other words, their data is flexible—which in turn means their business can be nimble and innovative. The thing to remember is, simply amassing data doesn't lead to innovation because data alone is not insight. What does convert data into insight is access—easy, quick, built-in access for its intended users. Without access, there is no insight. Without insight, there can be no innovation. It's that simple. And also, that complex.

Here's the kicker though. Flexibility is a direct result of controlled costs. If moving and accessing data sends your budget into a frenzy of API charges and egress fees then flexibility goes out the door, taking innovation with it.
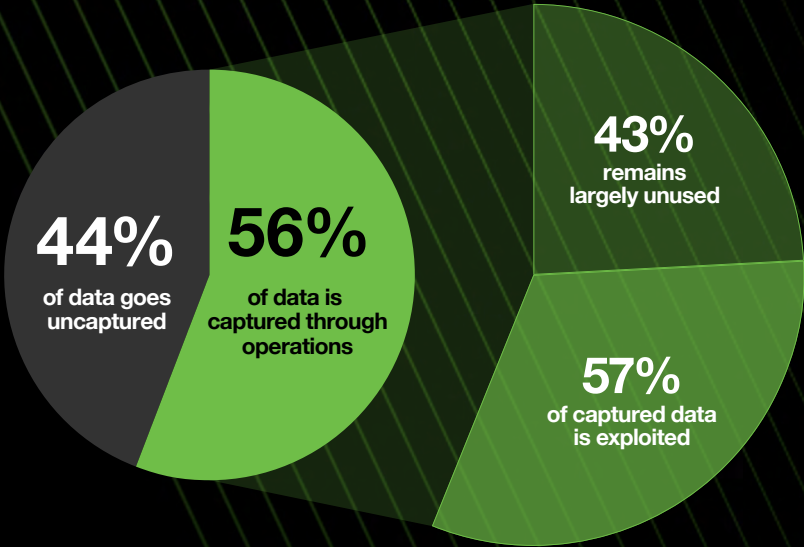
Third: Store More Data. Derive Its Value Faster.

Today's economy is driven by data, meaning companies that have the most data and activate it the fastest are the ones winning. As we've discussed, cost and flexibility play a huge role in your ability to do this.

Remember when we mentioned 220ZB of predicted data growth by 2026? Once that hits, some businesses will be prepared. Others won't be. Those that are will have infrastructure in place that allows them to efficiently and affordably capture, store, move, and manage loads upon loads of data. For these businesses, a massive figure like 220ZB is a boon because they'll be able to convert the large majority of their data into business value. This in turn keeps them competitive. However, for those who don't have the right infrastructure in place, this level of exponential data growth promises to be a burden, and not a small one. If current trends are any indication of the future, then the unfortunate truth is that most businesses will fall into the latter category.

Here's what the trends say: Today, 44% of enterprise data goes uncaptured. Of the 56% that is captured, 43% goes unused. This means that only 32% of available data is put to work. The remaining 68% goes unleveraged. In other words, there's an absurd level of loss taking place when it comes to business data and its potential value.

Only 32% of enterprise data is put to work.

44% of data goes uncaptured

56% of data is captured through operations

43% remains largely unused

57% of captured data is exploited

Now, take these numbers and factor in 220ZB of massive data explosion, as well as the increasing need for a (currently broken) multicloud ecosystem to handle this explosion, and things start to look like one of those disaster movies—where a tidal wave of data is headed this way and there sits your business in a rickety boat, frantically paddling to no avail.

At Seagate, this is where we yell, "CUT!" and declare that the scene sucked and it's time for another take. Why? Because this isn't how the future should play out. In this scenario, innovation and business success are limited. As a result, so is human potential.

Constraints Lead to Innovation

Over the past 40+ years, we've talked with thousands upon thousands of business leaders across the globe from every industry you could imagine. We've listened and learned about the constraints they face and their data disaster stories, and we've used these findings to inspire and shape solutions that help turn the tides. Seagate CEO Dave Mosely summed this up perfectly when he said, "Innovation is not driven by trends, but by the need to create more value under constraints—and there are constraints everywhere. Access to the full value of data is one of them."

The messiness of today's multicloud is just that—a constraint. One you can, in fact, overcome. We'll tell you how in a moment, but first, let us tell you a story.

Overcoming Our Own Multicloud Misery

Here at Seagate, we've battled our own fair share of data-related headaches in the multicloud, namely around TCO and security. Other hurdles have included a lack of long-term data recovery options, confusing pricing models from backup storage vendors, and of course, managing data sets that grow bigger by the day.

As users of the multicloud, what we've experienced is that most public clouds don't support big data analytics from an economics perspective. Thanks to complicated storage tiers and additional charges for moving and accessing data, we've thrown away more of ours than we'd like to admit—all because it was too expensive to store. The irony of course is that our business's sole focus is to empower people's data through storage. Why couldn't the multicloud return the favor to us? Irritated by a lack of solutions and anxious over our own data future, we were left with one option: build a solution ourselves.
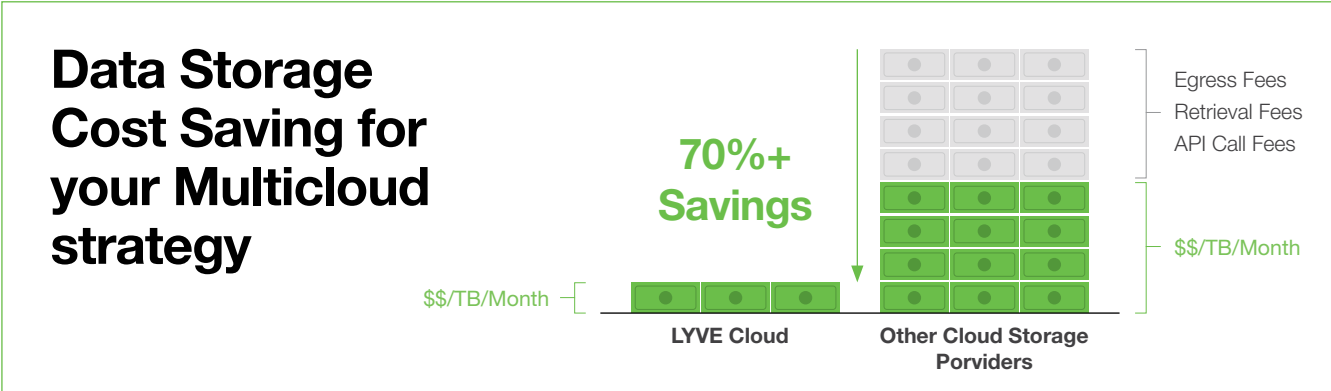
Companies that minimize data costs and maximize data innovation are the ones that win.

Seagate Lyve Cloud: Multicloud Freedom as a Service

Simple, trusted, and efficient object storage designed for multicloud—that's the solution we crafted and it's called Lyve™ Cloud. Today we offer this object storage as a service (STaas) to other businesses who, like us, rely on the multicloud and need it to work the way it was intended. Where you can control costs, access data anytime, move it around however you want, and never have to delete it. Whatever your existing multicloud storage strategy is, the Lyve Cloud platform is designed to integrate seamlessly. Here's what else you can expect:

- **Multicloud Freedom:** No egress fees, no API charges. Just complete freedom to transfer data seamlessly across public and private cloud environments.

- **Predictable TCO:** Pay only for the storage capacity you use, that's it. Simplicity like this means lowering your cloud bill by up to 70%.

Data Storage Cost Saving for your Multicloud strategy

70%+ Savings

Egress Fees
Retrieval Fees
API Call Fees

$$/TB/Month

$$/TB/Month

LYVE Cloud   Other Cloud Storage Porviders

- **Always-On Availability:** Lyve Cloud makes data available. Access it anytime without costly delays.

- **World-Class Security:** Safeguard your data assets with the most stringent, globally recognized security standards.

- **High Data Resiliency:** Prevent data deletion, manipulation, and data integrity challenges with highly durable and immutable object storage that supports data protection and anti-ransomware strategies.

- **Frictionless Object Storage Deployment:** Get up and running fast with user-friendly, S3-compatible application integration.

- **Tape Migration and Storage Services:** Seamlessly migrate backup tapes, digitize films, and move tape archives to the cloud.

- **Validated Partner Solutions:** Seagate works with industry-leading providers to ensure seamless compatibility between Lyve Cloud and partner solutions and the best possible user experience.

- **Demonstrated Expertise:** Our HDDs have long been trusted by cloud experts like AWS, Azure, and Google. Now you can get storage services directly from us in a cloud we've precision engineered for complete efficiency, total security, and improved TCO.

Ways to Use Lyve Cloud

Aside from seamless integration and capacity-based pricing, let's dig into the ways Lyve Cloud benefits specific use cases.

Data Analytics

For data that's generated by videos, applications, IoT devices, and more, Lyve Cloud helps you unlock its value with your choice of compute for S3 workloads. Reliable, secure, and highly available, Lyve Cloud object storage enables frequent access to large data sets without cost-prohibitive extras—bringing predictable economics to big data analytics, MLOps, data engineering, and AI.

Backup

Losing data to human errors, disaster events, or ransomware can destroy entire businesses. Don't let yours be one of them. With Lyve Cloud, you can fortify your 3-2-1-1 backup strategy with always-on cloud storage. Multi-region availability, replication, and object immutability make this storage-only cloud an ideal repository for backup storage and disaster recovery.
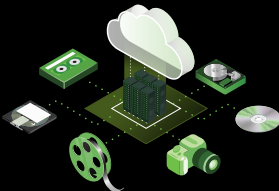
Media and Entertainment

Is your video and image data being siloed by inefficient infrastructure? Deploy easy-to-use cloud storage for increasingly higher resolution media and growing stores of content while paying only for the capacity used. This scale-out repository integrates seamlessly with post-production applications and AI-enabled technologies so you can easily access data for activation and reuse. Improve media production workflows and remote collaboration for distributed media teams with always-on availability and object immutability for ransomware protection.

Video Management

From manufacturing lines and security environments to traffic control systems, AV training data, and beyond, Lyve Cloud lets you access live and recorded video data from anywhere on demand and retain it securely and efficiently for as long as you need with a hybrid or hosted video surveillance as a service solution (VSaaS) or video management system (VMS). End-to-end user access control and ransomware protection safeguard evidentiary data while seamless connection to S3-compatible solutions supports guaranteed retrieval of critical data while enabling analytics that drive business value.

Conclusion

The number of people living on this planet has surpassed 7.8 billion, with more and more of us working from home and engaging with technologies centered around IoT, edge computing, and AI. Data proliferation (we'll say it again, 220ZB by 2026) is happening whether any of us are ready or not.

Mining and analyzing this data can open doors to a new wave of innovation, higher operational efficiency, and new sources of revenue. In order to harness this much data and garner this much value, we'll absolutely need the multicloud, and we'll need it to be optimized. Not messy. Not costly. Optimized.

That's why a STaaS solution like Lyve Cloud should be central to any multicloud strategy. It simplifies the complex cloud management system; it enables organizations to store data they might have historically been forced to delete; and as a result, business leaders can cost-effectively make the most of their data—spurring better outcomes and greater opportunities for growth.

Ready to Learn More?

Talk to an expert **www.seagate.com**

References:

https://www.seagate.com/our-story/rethink-data/

https://insidebigdata.com/white-paper/new-study-details-importance-of-tco-for-hpc-storage-buyers/

Resources:

https://www.seagate.com/services/cloud/storage/

seagate.com

SEAGATE