



SEAGATE

WHITE PAPER

# THE MULTICLOUD MANIFESTO

Managing Data Complexity at Scale

# CONTENTS

- 3** INTRODUCTION
- 4** SETTING THE STAGE:  
CLOUD COMPARISONS
- 6** FINDING THE MIDDLE GROUND
- 8** EXPLORING THE  
MULTIPURPOSE ADVANTAGE
- 10** EXAMINING THE EVOLUTION  
OF CLOUD-CENTRIC ARCHITECTURE
- 11** ASSESSING THE ROLE OF  
CLOUD SERVICES AT SCALE
- 12** MANAGING MULTICLOUD COMPLEXITY
- 13** CONSIDERING A  
THREE-STEP STRATEGY
- 14** CONCLUSION  
MAKING THE MOST OF MULTICLOUD



## Introduction

Every enterprise today is part of a data economy. In this economy, data is currency; it's how value is created and exchanged. Data is the new natural resource—the fuel of innovation.

An enterprise data economy is a global digital ecosystem in which data is gathered, organized, and exchanged by a network of cross-functional enterprise departments for the purpose of deriving value from the accumulated information.

Data collected and managed in an enterprise data economy must be securely stored in dedicated storage repositories, which are physically distributed across a mix of on- and off-premises data center environments—including traditional on-prem and co-location data centers, private cloud, public cloud, hybrid cloud, and ultimately, *multicloud*.

Multicloud describes an enterprise data infrastructure that uses more than one public cloud and increasingly includes private cloud components. In practice, multicloud can mean that an enterprise houses some data in public clouds and some in on-prem private cloud and may move data between them. The enterprise's data across all cloud elements is orchestrated through data management.

The multicloud landscape, while being exceedingly complex, offers great potential if data managers can break down and understand how to exploit the advantages of its various constituent capabilities.

The goal of a multicloud data architecture is to yield optimal data economics; that is, to create the most efficient infrastructure possible so as to capture the growing data opportunity, while also moving data to the right place within the organization so it can be converted to insight and drive competitive advantage.

A flexible infrastructure approach to multicloud will enable data managers to choose from owned, leased, or shared space—or a combination of all three. It's about whatever is needed to seamlessly scale to meet business requirements.

This white paper focuses on the emergence of multicloud data management strategies as a dominant paradigm for organizations with large and increasing storage demands, as well as the need to control the flow and use of data more closely than traditional cloud enables.



## Part 1 – Setting the Stage: Cloud Comparisons

Cloud solutions are typically defined by type: public, private, hybrid, and multicloud.

Public solutions offer multitenant compute and storage solutions off premises and on demand. Services such as Amazon's AWS and Microsoft Azure fall into this category.

Private cloud deployments are custom-built, single-tenant offerings that provide greater control over applications, resource utilization, and security.

Hybrid architectures use a mix of public and private clouds to maximize organizational efficiency by combining key benefits, such as public cloud load balancing and secure private storage.

Multicloud options use two or more clouds of any type. Although similar to hybrid deployments, there's a significant difference. Hybrid deployments effectively operate in series, with public and private services intermingling to produce specific results, while multicloud solutions run in parallel, with each cloud service running independently.

Each cloud approach features unique strengths and weaknesses:

### Public cloud

- Excels at: Shifting capital costs to operational expenses by moving resource management and security off site. In a public cloud framework, for example, providers are responsible for maintaining server hardware, completing software updates, and remediating security issues.
- Struggles with: Storing and accessing massive data sets. As data volumes increase, so does the time required to move information back and forth from on-site to in-cloud servers.

### Private cloud

- Excels at: Combining public and private processes to improve operational output. For example, hybrid solutions are often used to help enhance performance by combining public application services with on-premises data storage.
- Struggles with: Rapid scalability to meet the emerging demands of data-driven analytics that inform strategic decision-making. Private clouds also struggle to balance price and peak demand. While over-provisioning ensures peak demand performance, unused resources are costly for typical traffic throughput. Under-provisioning reduces cost, but also puts companies at risk of getting overwhelmed.

### Hybrid cloud


- Excels at: Combining public and private processes to improve operational output. For example, hybrid solutions are often used to help enhance performance by combining public application services with on-premises data storage.



- Struggles with: Reliably getting the best of both worlds. Evolving complexity, seamless data placement, and orchestration concerns often conspire to frustrate hybrid outcomes.

**Multicloud**

- Excels at: Data management at scale by combining mass-storage clouds with centralized-compute clouds to deliver better data economics.
- Struggles with: Increased complexity as IT teams must manage the ongoing interaction of multiple clouds simultaneously.




## Part 2—Finding the Middle Ground

Enterprises are embracing the cloud to gain granular resource control, streamline IT spending, and deliver on big data potential. It makes sense: With data volumes rapidly increasing, 68% of available information already goes unused, according to research from the International Data Corporation cited in Seagate’s recent *Rethink Data* report. And with enterprise data on track for 42.2% annual growth, organizations require solutions that can leverage this resource at scale.


But as cloud shifts accelerate, the limitations of single-provider models begin to emerge. Early adopter benefits offered by the public cloud—such as pay as you go, scale as you need, and faster time to market—have given way to challenges that include latency, cost control, and privacy. Private clouds, meanwhile, are often hampered by existing infrastructure limits and increasing management costs.

Multicloud offers a middle ground. The concept is simple: Enterprises use two or more cloud providers to meet specific operational goals and deliver reliable return on investment (ROI). Recent data cited in the Rethink Data report shows that multicloud is on the minds of IT leaders: 37% have already deployed a multicloud strategy across their IT infrastructures.

In practice, however, multicloud deployments come with their own challenge: complexity. In fact, 61% of enterprises point to the challenges of managing data in a multicloud environment over the next two years. With data now widely and evenly distributed across enterprise, on-premises, cloud, and edge computing repositories, enterprises must manage multiple workflow streams and disparate management tools simultaneously to reduce conflict across cloud service architectures.



As a result, finding a cloud middle ground is only the first step in enterprise success. By developing a mandate that defines both key roles and critical limitations for public and private cloud solutions, organizations can effectively leverage the scope of multicloud benefits and manage complexity at scale.



## Part 3—Exploring the Multipurpose Advantage

Business success is a trifecta of revenue, profit, and customer satisfaction. To achieve all three, enterprises are moving away from one-size-fits-all offerings to best-in-class solutions that align with unique organizational challenges. This shift to purpose-driven cloud usage aligns with the democratization of cloud services at scale. Adoption is no longer restricted by public or private cloud classifications; function now forms the defining characteristic.

The result is a multipurpose menagerie of services that can quickly outpace an IT department's ability to effectively operate and optimize. Before making the move to multicloud, companies should evaluate current proficiency across three key categories:

- *Management.* The ability to effectively manage data storage, movement, and security across cloud services is critical to multipurpose success. With 37% of enterprises already moving toward multicloud environments ([per the Rethink Data report](#)), disparate data management skills are essential.
- *Capacity.* Data capacity needs are quickly evolving as companies combine public and private services to deliver optimal outcomes. As a result, organizations need capacity strategies that account for current needs, anticipate emerging demands, and can effectively manage the divide between on-premises storage requirements and public cloud capacities.
- *Efficiency.* Agility is a hallmark of the cloud but is often lacking in enterprise IT strategies. While piecemeal adoption of new solutions and services to meet specific needs offers in-situ benefits, businesses must address their data efficiency at scale to identify the best approach to cloud service expansion. According to [recent Deloitte data](#), almost half of all companies surveyed said cloud adoption is more complex than they had expected. As a result, building efficiency is critical.

In addition to evaluating current competencies, companies must also take steps to limit sprawl, improve efficiency, and control spending. Critical considerations include:


- *TCO calculations.* Differing cloud frameworks demand different total cost of ownership (TCO) calculations. If budgetary constraints inform current IT decision-making, public cloud solutions offer the lowest TCO-to-resource ratio. Private clouds demand more up-front TCO but may outpace pay-as-you-go public clouds over time.
- *Differing data functions.* Purpose defines priority. Sensitive, compliance-controlled data, such as personally identifiable information or financial records, is often best kept under the auspices of private cloud solutions, while anonymized data can deliver significant value when leveraged by large-scale analysis tools powered by public clouds.
- *Increasing capacity requirements.* Physical hardware can't always keep pace with increasing capacity requirements. In this case, public cloud storage offerings can help bridge the gap between needed space and existing storage. At the same time, as data centers continue to evolve toward more open and composable architectures, and as open source software





solutions reduce the costs of data orchestration, IT budgets for private cloud and on-prem data centers can be expected to more readily absorb more storage costs than in the past—enabling rapid deployment of new storage systems to meet requirements.

- *Movement and repatriation of data.* As data analytics tools evolve to include artificial intelligence and machine learning capabilities, many companies are repatriating their data—moving it back from public clouds to private stacks—to enable on-site, on-demand analysis for actionable results.
- *Evolving data security demands.* Both consumers and compliance organizations demand due diligence from first-party data owners. As a result, diligence informs deployment. While public cloud security has increased significantly over the past five years, ultimate responsibility can't be shifted to providers.



## Part 4—Examining the Evolution of Cloud-Centric Architecture

Beyond basic considerations and specific use cases, the evolving nature of cloud solutions speaks to adoption at scale.

Consider data storage: as data volumes increase exponentially, file-based storage now faces a challenge from new solutions, such as object-based frameworks, which empower the transfer of massive data sets at speed.

Storage architecture is also evolving as options, such as hyperconverged and disaggregated, composable solutions, emerge. While hyperconverged offerings leverage software-defined clustering to provision storage resources on demand, disaggregated, composable systems decouple storage, network, and compute processes from physical hardware to create a shared pool of resources that can be individually assigned or used in concert.

Open source initiatives, meanwhile, are changing the fundamental nature of the cloud. From open source software developments that allow organizations to customize key functions and frameworks to open source architecture efforts—such as storage devices that are programmable and modifiable—the shift away from proprietary provisioning will impact long-term cloud decision-making.

Organizations also must account for the movement of data and compute tasks away from central cloud services to edge infrastructure. Informed by increasingly complex processes handled by connected, intelligent devices at the point of origin—rather than being shunted to public or private clouds for analysis—edge computing modifies the value proposition of both public and private cloud services, as well as their impact on operational outcomes. As noted by **The Enterprisers Project**, 30% of IT budgets are now earmarked for edge computing, meaning organizations must be prepared to manage functions at the fringe.



## Part 5—Assessing the Role of Cloud Services at Scale

With the scalability of public cloud deployments, the specificity of private services, the continuity of hybrid options, and the opportunity presented by multicloud solutions, companies are now spoiled with choices.

But as organizations add new storage and compute services across IT environments at scale, finding best-fit solutions becomes a challenge. What seems like a cut-and-dried case for public solutions may suddenly become an ideal private initiative when composable architecture is introduced or object storage is supported.

Here, an abundance of choice can lead to cloud decision fatigue—driving the adoption of type-based services without a reliable road map. The result is a multicloud environment bereft of organizational oversight that naturally lends itself to compute complexity and operational inefficiency.

To deliver on purpose-driven potential, role-specific assessment is essential across each cloud framework:

- *Public clouds.* Public solutions offer the lowest barrier to entry for enterprises that are shifting to cloud services. As a result, they're often on the front lines of compute conversion. By moving key processes into the cloud, organizations can free up much-needed overhead in local environments.

But, as noted by Seagate senior vice president and CIO Ravi Naik in a recent [Venture Beat](#) piece, what worked at 1 petabyte (PB) in the public cloud won't work at 100PB. As data volumes increase, issues are emerging around billing, transparency, and cost predictability.

- *Private clouds.* Private clouds offer control. As a result, they're the go-to option for organizations looking to embrace resource scaling without compromising internal security.

With control comes potential compromise. From insider threats to outdated virtual machine (VM) images, if private clouds aren't properly managed, staff accidentally misuse internal cloud services, or IT teams forget to update VM security software and settings, critical data could be compromised.

- *Hybrid clouds.* Hybrid clouds are all about balance. By combining public and private resources, businesses can both streamline operations and simplify resource procurement at scale, effectively finding a middle ground between the substantial resources available to public cloud users and the significant control offered by private cloud alternatives.

When it comes to compliance, however, hybrid offerings often face challenges. With public and private services intertwined, the line between on- and off-site storage can begin to blur if sophisticated data orchestration tools with effective policy-based data management are not in place, putting organizations at risk of accidentally exposing critical data by association rather than intention.

- *Multiclouds.* The multipurpose nature of evolving cloud services has created a natural trend toward multicloud deployments. By leveraging cloud services in parallel to produce specific outcomes, organizations achieve a measure of the cloud's theoretical intent. In practice, this means decoupling provider and process to deliver on-demand results.

Where multicloud often meets its match, however, is complexity. While hybrid clouds include overlapping edges, multiclouds come with the need to create inter-cloud cooperation at scale.



## Part 6—Managing Multicloud Complexity

As companies broaden their cloud horizons with the adoption of public, private, and hybrid solutions, multicloud becomes the de facto structure of cloud environments.

The shift to multicloud makes sense as businesses look to leverage specific functions rather than develop vague organizational frameworks. This multicloud move delivers the best of both worlds: scalable public resources for large-scale storage or data analytics, paired with private cloud solutions to handle mission-critical tasks.

Despite this alignment of purpose and process, however, the increasing scope of multicloud solutions creates its own challenge: complexity.

- *Separate workflows.* While using distinct cloud processes can reduce user frustration by providing purpose-built tools, deploying an application-based-outcome approach means each process leverages its own separate workflow. Without effective monitoring and management across cloud networks, these processes begin to conflate and conflict, in turn reducing overall efficacy.
- *Disparate management tools.* Many cloud solutions come with their own set of management tools for handling resource deployment, permissions, and application calls. But these tools often lack interoperability, forcing IT teams to spend time on redundant—albeit slightly different—processes across multiple, similar cloud applications.
- *Lack of unified security.* Multicloud solutions naturally introduce potential security concerns. Since each cloud-based application interacts with network security solutions to a different degree and requires different resource permissions, fragmented security is often the result.
- *Data movement and sharing.* Data is handled and stored differently by public and private cloud applications, in turn creating operational roadblocks when users attempt to export or import key data sets for analysis or reporting.
- *Making collected data usable.* Thirty-nine percent of enterprises now highlight data usability as their biggest challenge. While multicloud environments empower data collection at scale, they also introduce complexity as businesses look to manage multiple data sources simultaneously.



## Part 7—Considering a Three-Step Strategy

The multicloud is gaining ground as organizations embrace the application-driven outcomes delivered by purpose-built cloud solutions at scale. Where these solutions exist—on public servers or private stacks—is no longer their defining characteristic.

But deployment alone isn't enough to deliver on the potential of big data. To make the most of multicloud movement, enterprises need a three-step strategy:

1. Identify data architectures

Before making the move to multicloud, it's critical for companies to identify current data architectures—including where data is stored, how it is used, and where processes could be improved—to define desired operational outcomes. With specific goals in mind, enterprises are better prepared to make multicloud decisions that deliver line-of-business benefits.

2. Assemble cloud resources

Purpose-built clouds are now commonplace across both public and private deployments thanks to emerging solutions, such as customizable, composable infrastructure. As a result, it's critical for companies to assemble multicloud network, compute, memory, and storage functions that are capable of addressing current needs and meeting future challenges.

3. Enable seamless orchestration

Multicloud deployments only deliver on performance and process mandates when cloud elements operate seamlessly. Realizing ROI on this move to multicloud means building on defined architectures and assembled resources with robust data orchestration and management solutions capable of reducing friction across key cloud functions.



## Conclusion

The shift to cloud is inevitable as data volume, variety, and velocity increases exponentially—and the demand for actionable insights based on intelligent analysis of new data sets is keeping pace. To achieve operational outcomes at both speed and scale, companies are moving beyond dichotomies of public and private to embrace hybrid frameworks and multicloud models that deliver specific, application-based outcomes.

But the expanding impact of multicloud environments introduces a new issue: complexity. Tasked with managing multiple clouds across storage, compute, security, access, and analytics applications, IT teams are faced with a growing service disconnect. When every cloud process comes with its own workload structures, management controls, and security settings, enterprises must conquer the challenge of creating constructive communication across disparate delivery channels.

Multicloud solutions decouple key cloud functions from familiar cloud forms. With the right strategy, enterprises can both reduce complexity and facilitate multicloud management at scale.

Ready to Learn More?

Visit us at [seagate.com](https://www.seagate.com)



### Seagate Private Cloud Solutions

[www.seagate.com/solutions/cloud/private-cloud/](https://www.seagate.com/solutions/cloud/private-cloud/)

### Lyve Rack – Open affordable, and durable object storage solution

[www.seagate.com/products/storage/object-storage-solutions/lyve-drive-rack/](https://www.seagate.com/products/storage/object-storage-solutions/lyve-drive-rack/)

### CORTX – Open source mass capacity object storage

[www.seagate.com/products/storage/object-storage-software/](https://www.seagate.com/products/storage/object-storage-software/)

© 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. 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. TP720.1-2101US, January 2020

