

WHICH CLOUD STORAGE SERVICE DELIVERS THE
PERFORMANCE YOU NEED? COMPARING
IBM CLOUD OBJECT STORAGE AND AMAZON S3

Stratecast

F R O S T & S U L L I V A N

An Executive Brief Sponsored by IBM

Lynda Stadtmueller
Vice President – Cloud Services

July 2016

INTRODUCTION

Not long ago, many organizations treated data storage much like physical storage: they just tucked it away in any available space. Now, data is recognized as a valuable corporate asset that can yield business-impacting intelligence. To help make decisions and run operations, companies are collecting and using exponentially more data—both structured data, which is growing linearly, and unstructured data, which is growing exponentially.

To manage the volume, variety, and velocity of data, enterprise storage and IT executives have adopted infrastructures based on object storage—a scalable storage structure in which data and its metadata are separated and stored in an “object format.” Rather than a rigid hierarchical structure of a file system, the objects are stored in a flat name space, with applications accessing diverse data objects based on the object name or identification. Such storage systems enable enterprises to scale and tune storage attributes to precisely meet their application needs.

The growing interest in object storage parallels the proliferation of unstructured data—for example, video, photos, emails, messaging, documents—and the increase in cloud and mobile applications that produce and access such data. In the past, what most businesses considered “usable” data could be captured in a structured format of rows, columns, and fields—for example, financial or customer data. But today, all that has changed, as enterprises embrace trends such as cloud, mobile, social business, analytics, and Internet of Things. Businesses are now amassing and utilizing escalating volumes of unstructured content. In addition, developers of mobile and cloud applications require storage that is Web-scale and Application Programming Interface (API)-friendly. Object storage has shown potential as a solution that can meet evolving needs.

Many enterprises will struggle in their search for an object storage solution that meets their criteria for flexibility, performance, scalability, control, transparency, and cost-effectiveness, especially as they look to the cloud. Many cloud object storage services, including Amazon S3, deliver capacity with little control, meaning that enterprises have little ability for the enterprise to adjust performance parameters to meet the needs of their applications.

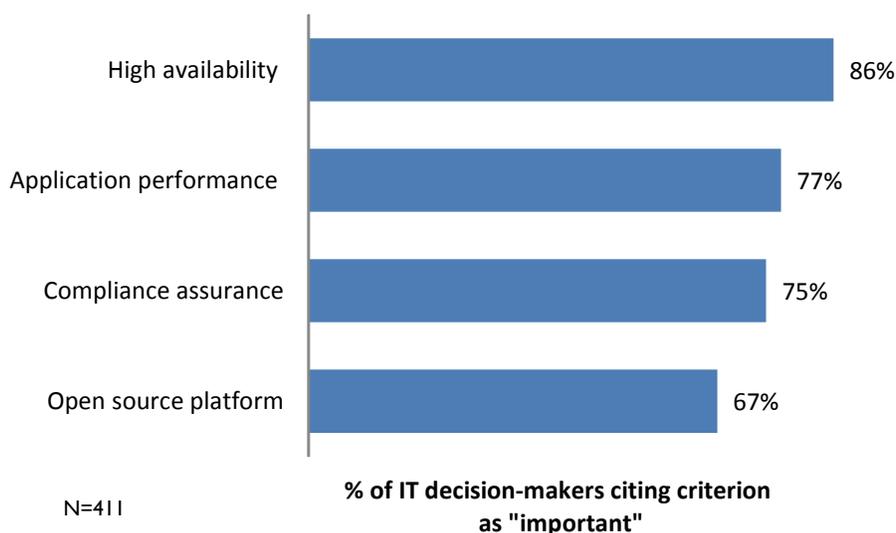
In contrast, IBM Cloud Object Storage offers users a high-degree of flexibility and control over their object storage, thanks to the feature-rich IBM technology platform. Furthermore, IBM Cloud Object Storage is the first truly hybrid cloud storage system with functionality delivered both in the IBM Cloud data centers around the globe and for on-premises deployment.

In this paper, we provide a head-to-head comparison of IBM Cloud Object Storage and Amazon Web Service (AWS) Amazon S3; looking at features, functionality, and—based on comprehensive tests conducted by IBM—performance.

THE ELUSIVE PRIVATE STORAGE CLOUD

As enterprises build out their hybrid IT strategies, they are increasingly choosing private or dedicated cloud services to run their workloads. IT decision-makers surveyed by Stratecast define “hosted private cloud” as single-tenant infrastructure that is not shared by other businesses. Those survey respondents expect a lot from their hosted private clouds, as shown in Figure 1.

Figure 1: Most Important Criteria in Selecting a Hosted Private Cloud Service



Source: Stratecast

However, few cloud service providers offer a private or single-tenant option for their cloud storage services. For example, AWS offers enterprises the option for single-tenant compute instances; however, the company’s object storage is multi-tenant, hosted on shared storage infrastructure.

The lack of a single-tenant storage option can pose a challenge. Multi-tenant architecture can negatively impact workload availability and performance, because the resident workloads contend for infrastructure resources. Organizations’ compliance requirements may prohibit placement of critical data on shared devices. Furthermore, an open architecture is desirable to provide the flexibility enterprises require to move and access storage. Without the option of a true open and private storage cloud, enterprises’ storage needs will be unmet.

THE CHALLENGE OF HYBRID CLOUD STORAGE

If a true private storage cloud is difficult to find, enterprises may struggle even more to implement a hybrid storage cloud. According to the Stratecast survey, enterprises define “hybrid” as an integrated environment comprising both on-premises and hosted cloud infrastructure. Yet, most object storage solutions available on the market today are built for either cloud or on-premises deployment—not both.

This presents a challenge to businesses that operate in a hybrid environment with hybrid business processes, in which the applications and the stored data they access may be dispersed or split across multiple environments. Consider, for example, a scenario in which an industrial organization builds a Web-facing partner portal to provide its network of installers and maintenance technicians with mobile access to a comprehensive library of schematics. The portal is deployed on the cloud; the proprietary schematics are stored in an on-premises object storage system. Now, suppose the organization decides to enhance the effectiveness of the library by including video tutorials, which are stored in the cloud to take advantage of the scalability and on-demand usage. The result is a hybrid configuration that must deliver a seamless, consistent experience to users regardless of where the app and content are deployed.

A unified hybrid storage environment allows businesses to place storage where it makes most sense today, and to make a different decision tomorrow, as business needs change. For example:

- A business may choose to align storage with the deployment model that hosts the application that primarily accesses the data—placing storage for cloud apps in the cloud, and storage for on-premises apps on the premises.
- A business may look to move aging storage across environments. As an example, a hospital may choose to store medical images on-premises, where they can be easily accessed during the course of treatment; but to archive the images in a secure and compliant private cloud after a period of time, to leverage the cloud’s scalability and cost-efficiency.
- To manage an Internet of Things workload, a business may implement a two-tiered storage structure, with local, on-premises devices collecting, storing, and conducting preliminary analysis on sensor data; but transmitting metadata to the cloud for aggregation, deeper analysis, and permanent storage.

In a hybrid environment, regardless of where the storage is hosted, enterprises expect to treat the various environments as pools of capacity, with visibility and control across environments, and the ability to migrate storage data with minimal friction. This requires a unified storage platform that extends across both the premises and cloud.

A NEW TYPE OF CLOUD OBJECT STORAGE FOR YOUR HYBRID CLOUD: IBM CLOUD OBJECT STORAGE

Until recently, enterprises did not have an option to deploy a high-performance object storage solution across both the cloud and the on-premises data center. That changed with the introduction of IBM Cloud Object Storage, which is designed to meet the needs of businesses that expect a lot from their valuable data. The series of offerings provide higher levels of flexibility, control and functionality than competitive services such as Amazon S3.

IBM Cloud Object Storage (built up from the proven technology that IBM acquired through the acquisition of Cleversafe, Inc. in 2015) enables businesses to take control of their cloud, and select the optimal deployment model, costs, and performance for each workload. Key features include:

Feature	IBM Cloud Object Storage	Amazon S3
Single/multi-tenancy options	<ul style="list-style-type: none"> ▪ Multi-tenant ▪ Single-tenant (Dedicated Service) 	<ul style="list-style-type: none"> ▪ Multi-tenant only
Deployment options	<ul style="list-style-type: none"> ▪ On-premise (appliance or licensed software) – managed by enterprise or IBM ▪ IBM Cloud ▪ Unified hybrid deployments 	<ul style="list-style-type: none"> ▪ Cloud only
Customization and control	<ul style="list-style-type: none"> ▪ With Dedicated Service, dynamic control over performance, on a workload basis ▪ Visibility and reporting 	<ul style="list-style-type: none"> ▪ No customized control or workload visibility ▪ Standard reporting
API support	<ul style="list-style-type: none"> ▪ Openstack Swift ▪ S3 Compatible API ▪ Simple object API ▪ NSF/SMB 	<ul style="list-style-type: none"> ▪ S3 only

COMPARING PERFORMANCE: IBM CLOUD OBJECT STORAGE VERSUS AMAZON S3

Performance is top of mind for enterprises considering cloud storage providers. Fully 80% of enterprises surveyed by Stratecast cite “end-to-end application performance assurances” as their top criterion for selecting a cloud provider. Conversely, 56% of businesses list concerns about “poor or inconsistent application performance” as a reason not to use cloud.

However, performance of cloud storage (measured in latency and consistency) varies widely among providers. Providers’ technology, deployment options, and tools contribute to the differences, which can impact your businesses’ productivity, agility, flexibility, and overall costs.

But differences in performance are not easy to discern from providers' Web sites, which tend to focus on price and capacity. Unfortunately, many enterprises choose a cloud object storage provider largely based on price-per-Gb, without realizing that performance may suffer.

To test the point, IBM ran identical object storage workloads against IBM Cloud Object Storage and Amazon S3, over a period of time. Stratecast reviewed and verified the methodology, assumptions, results, and analysis. The test parameters and results follow.

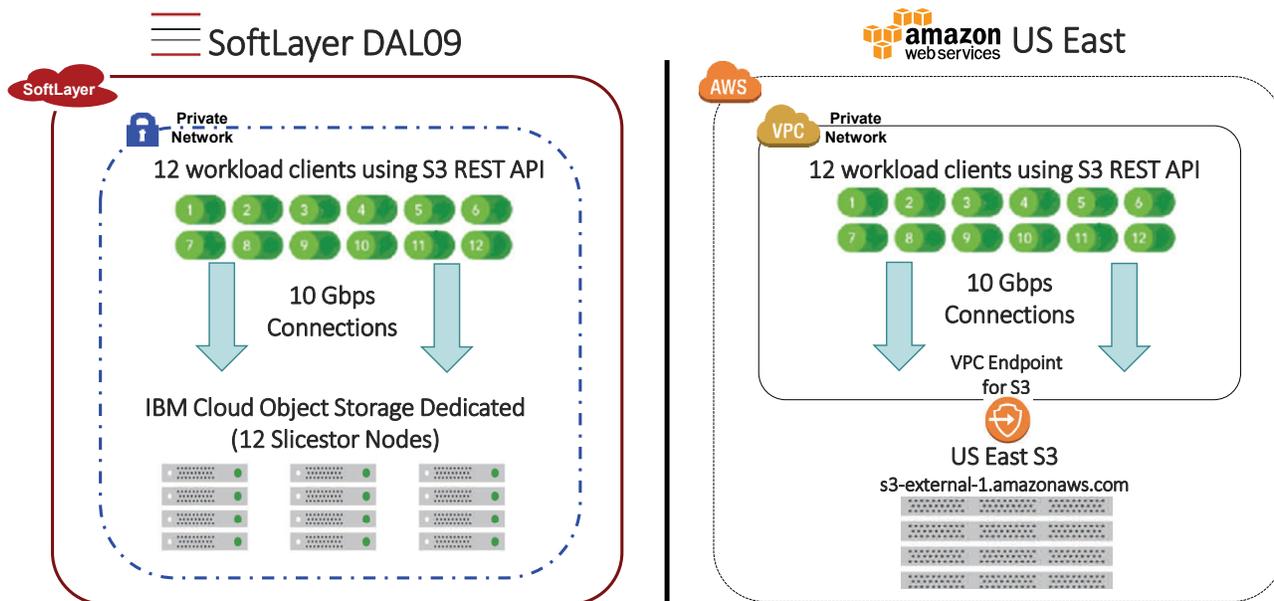
Test Objective: Measure maximum “read” and “write” performance (in operations per second), throughput (maximum amount of data moved per second in megabytes per second) and transaction latency (in milliseconds per operation) over time, under comparable and optimal configurations.

Workload Specifications: Proprietary object storage workload; 12 workload clients using S3 REST API, 10 Gbps connections. Storage reflected a range of object types (including metadata, email, Web page, audio, video, and images) and sizes (ranging from 10KB to 100 MB).

Configurations: In the test, the optimal environment was selected for each workload.

- The IBM workload was run on IBM Cloud Object Storage Dedicated, the single-tenant option, using 12 storage nodes.
- Without a single-tenant option, the AWS workload was run on S3 in the US East data center.

Figure 2:
IBM Cloud Object Storage Dedicated Service and Amazon S3 Configurations



Source: IBM

Throughput Results

When the environments were configured comparably, **IBM Cloud Object Storage Dedicated Service delivered 1.9x higher “read” throughput and 3.3x higher “write” throughput than S3.** The higher performance of IBM Cloud Object Storage Dedicated can be attributed largely to the single-tenant architecture, which maximizes the server resources available to the workload.

However, “comparable” configurations actually means the “lowest common” configuration. While Amazon S3 offers no other configuration options, IBM Cloud Object Storage allows customers to adjust the configuration on a workload-by-workload basis to optimize for performance. For example, the customer can enable or disable features, such as object indexing, to speed performance for certain workloads. In contrast, S3 does not offer control or visibility into the internal operations of the service, so users are unable to monitor or modify features like indexing.

In the next phase of its test, IBM tested the same S3 configuration against an IBM Cloud Object Storage configuration that was optimized for “write” operations, by disabling the “indexing” function (a function that is not available with Amazon S3). **The write-optimized IBM Cloud Object Storage Dedicated delivered 1.7x faster “read” and 9.9x faster “write” performance than S3.**

Latency Results

In addition to throughput speed, another way to measure storage performance is by looking at latency (the length of time it takes to complete an operation). But there are different ways to look at latency, not all of them equally useful to understanding workload performance. Each operation takes a different amount of time to complete. Some users measure average latency across all transactions. However, the variation in complete times can be extreme, skewing the mean. For workloads in which consistency in performance is as important as the actual time to complete the operation, calculating a simple mean is not useful. A better approach is to measure latency by percentile, indicating what portion of the operations achieved a given latency result.

In this test, IBM configured the IBM and AWS tests comparably (with similar indexing configurations), using a constant request rate of 420 requests per second. Latency results were measured at the 95th percentile (indicating the latency incurred by that 95% of operations). **At the 95th percentile, IBM Cloud Object Storage Dedicated Service delivered latencies for “read” operations that were 3.5x-6.6x lower than Amazon S3, and “write” latencies that were 1.5x-11.5x lower than AWS.**

The dramatic differences in performance can once again be attributed to how the storage platforms are designed, as well as the availability of single-tenancy with the IBM Cloud Object Storage offer.

THE LAST WORD

When your business operations depend on consistent and rapid access to stored data, you need to be confident your storage will deliver. That's why the best way to compare among cloud object storage services is to look beyond flat price-per-GB, and instead focus on how the service will perform for a particular workload.

Unfortunately, too many cloud storage providers expect you to blindly trust their one-size-fits all configurations to deliver on your needs. Their inflexible cloud object storage services may deliver low-cost capacity, but offer little of the visibility, control, and flexibility that you need to optimize your storage on a workload-by-workload basis, whether in the cloud or in your own data center.

It's not just about handling accelerating data volumes and types; it's about driving business value from your storage. **Don't compromise: Look for the storage solution that delivers maximum flexibility, control, and performance across your hybrid environment.**

Lynda Stadtmueller

Vice President – Cloud Services

Stratecast | Frost & Sullivan

lstadtmueller@stratecast.com

To learn more about IBM Cloud Object Storage or the test parameters and results, see your account executive or visit <http://ibm.biz/cloudobjectstorage>.

Silicon Valley
331 E. Evelyn Ave., Suite 100
Mountain View, CA 94041
Tel 650.475.4500
Fax 650.475.1570

San Antonio
7550 West Interstate 10, Suite 400
San Antonio, Texas 78229-5616
Tel 210.348.1000
Fax 210.348.1003

London
4, Grosvenor Gardens,
London SW1W 0DH, UK
Tel 44(0)20 7730 3438
Fax 44(0)20 7730 3343

877.GoFrost • myfrost@frost.com
<http://www.frost.com>

ABOUT STRATECAST

Stratecast collaborates with our clients to reach smart business decisions in the rapidly evolving and hyper-competitive Information and Communications Technology markets. Leveraging a mix of action-oriented subscription research and customized consulting engagements, Stratecast delivers knowledge and perspective that is only attainable through years of real-world experience in an industry where customers are collaborators; today's partners are tomorrow's competitors; and agility and innovation are essential elements for success. Contact your Stratecast Account Executive to engage our experience to assist you in attaining your growth objectives.

ABOUT FROST & SULLIVAN

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies? Contact Us: Start the Discussion

For information regarding permission, write:

Frost & Sullivan
331 E. Evelyn Ave. Suite 100
Mountain View, CA 94041

Auckland

Bahrain

Bangkok

Beijing

Bengaluru

Buenos Aires

Cape Town

Chennai

Colombo

Delhi / NCR

Detroit

Dubai

Frankfurt

Iskander Malaysia/Johor Bahru

Istanbul

Jakarta

Kolkata

Kuala Lumpur

London

Manhattan

Miami

Milan

Moscow

Mumbai

Oxford

Paris

Rockville Centre

San Antonio

São Paulo

Sarasota

Seoul

Shanghai

Shenzhen

Silicon Valley

Singapore

Sophia Antipolis

Sydney

Taipei

Tel Aviv

Tokyo

Toronto

Warsaw

Washington, DC