

# TACO REPORT

## On-Prem Storage vs. Public Cloud Storage

Economic Advantages of On-Premises Object Storage  
vs. Public Cloud for Enterprise Data Storage

# Executive Summary

Today, organizations are challenged by a massive growth in the volume of enterprise data. It's an undisputed fact that today more data is being generated, processed, and stored than at any other time in history. And it's only going to increase from here. International Data Corporation (IDC) projects the "global datasphere" to grow from 33 zettabytes this year to a staggering 175ZB by 2025, representing a compound annual growth rate (CAGR) of 61%.<sup>1</sup>

Looking more closely, analysts at Gartner estimate that 80% of enterprise data today is "unstructured".<sup>2</sup> This is important because much of the new growth of data volume is unstructured data which includes email, video, voice recording, media files, life science and healthcare data, social media, IoT and sensor data, among others. It's also important because unstructured data is an ideal application for object storage. And with the rapidly-growing volumes of unstructured data volume, this is creating a corresponding increase in the need for object storage.

With this data growth putting a lot of pressure on organizations in every industry, technology leaders face a tough challenge: they must find a cost-effective way of storing and managing this growing amount of data without sacrificing performance, security, or service delivery.


As a result, many have turned to the public cloud. The flexible nature of cloud computing coupled with consumption-based pricing has made the public cloud a tempting option. But with time and experience — and enhancements to on-premises object storage solutions — many organizations are evaluating the public cloud in a new light. Some are turning to hybrid cloud approaches that employ both on-prem and cloud infrastructure, while others are repatriating applications and data to on-prem environments. According to a survey of IT managers by 451 Research, 58% of respondents said they are moving to a hybrid IT strategy.<sup>3</sup>

In this paper, we compare the economics of on-premises object storage with public cloud for enterprise storage. With demands on enterprise storage higher today than ever before, we first will compare the features and capabilities of on-premises object storage with public cloud and see how they stack up. Then, using a realistic customer scenario, we will analyze the total cost of ownership (TCO) to determine which approach is the most cost-effective for current as well as future enterprise storage needs.

## KEY TAKEAWAYS

- 5-year TCO reveals Cloudian on-premises storage is 65% less than public cloud.
- Cloud data access costs have a significant impact to overall TCO.
- On-premises object storage scales more cost-effectively than public cloud.
- Data durability for Cloudian on-premises object storage can be achieved at 11 nines and greater.
- On-premises storage enables faster data access and data restore with low-latency local storage.
- Comprehensive control of on-premises storage can provide improved security, privacy, and compliance vs. public cloud.

# Enterprise Storage Trends



For data-intensive applications, enterprise data storage is an increasingly critical part of the infrastructure —at a technical and operational level, and also in terms of associated efficiencies and costs.


Some of the application workloads driving unstructured data growth include: Healthcare, Media & Entertainment, Electronic Design Automation, Insurance, Education, Transportation, Oil & Gas, Retail, Analytics, and Life Science. As these and other workloads produce more and more data, there is an increasingly urgent need for cost-effective storage that does not sacrifice performance, security, or service delivery.

## Increased Demand on Enterprise Storage

These demands have resulted in organizations looking for new storage alternatives. And many of these organizations have already explored or are evaluating public cloud options.

The public cloud does offer several benefits. First of all, it's scalable —capacity can be added as needed to accommodate data growth. Second, customers do not need to manage hardware —support of the underlying infrastructure is the responsibility of the cloud provider. Third, public cloud provides an off-site location for data to protect against localized disasters.

*But is public cloud the best option for today's enterprise storage requirements? How does it compare to on-premises object storage?*



# On-Premises vs. Public Cloud Storage

In light of exponential data growth, heightened demand for functionality, stricter security requirements, and ever-tightening budgets, how does on-premises object storage compare to public cloud?

The following table offers a side-by-side comparison.

	ON-PREM OBJECT STORAGE	PUBLIC CLOUD STORAGE
<b>Data Access</b>	For local data, access is fast, <0.01 second.	Data is always remote. Access speed depends on location and WAN latency.
<b>Scalability</b>	Limitless growth within a single namespace.	Limitless growth within a single namespace.
<b>Efficiency</b>	Add capacity as needed. Scale-out architecture permits seamless growth. Cost per GB decreases as capacity grows.	Capacity can be added when needed. Cost per GB above 500 TB remains constant, regardless of growth.
<b>Searchability</b>	Data stored as objects. Metadata embedded within the object can be customized to facilitate data search	Same as on-prem object storage.
<b>Geographical Distribution</b>	Nodes can be deployed and shared across geographic regions.	Data can be stored in zones across different regions.
<b>Storage Management</b>	Policy-based. Fine-grained storage management policies. Unified view and control across all locations.	Storage infrastructure is managed by the cloud provider. Data is managed by the user.
<b>Data Recovery</b>	Self-healing media, integrated data protection. Data can be easily replicated between nodes and locations. Fast recovery over local network.	Data recovery time is dependent on available WAN bandwidth. Recovering 1TB over a 1Gb link requires 3 hours.
<b>Data Durability</b>	Erasur coding and replication options. Data distributed easily across multiple sites. Durability of 11+ nines, customizable as needed.	Standard offering of 11 nines.
<b>Compliance / Privacy</b>	Data resides on-prem within specified locations ensuring proper compliance via policy.	Data can reside in different zones. User must rely on cloud provider to manage.
<b>Control</b>	Complete control over data and underlying infrastructure.	No control over storage, only data. Cloud provider controls all underlying infrastructure.
<b>Security</b>	Managed as part of consistent enterprise security architecture framework.	Policy not always consistently applied leading to storage buckets and exposed data.
<b>Performance</b>	Fast ingest and operational performance. Being on-prem, latency impact is minimized.	Limited by the available WAN bandwidth.

## CASE STUDY

# Public Cloud Storage vs. Clouidian On-Prem Storage – 65% Savings

Comparing features and capabilities between on-premises object storage and public cloud is important, but it's only the first step. The second step is to define a use case with a realistic set of requirements and to configure each storage solution to satisfy those requirements. Then, analyze the two options.

For this analysis, we present a hypothetical customer scenario for an organization with one petabyte of storage capacity (or 1000 terabytes) for a five-year period. Of that total amount, we assume 10% (or 100 terabytes) of data is accessed each month. The comparison will be between storing the data in public cloud and Clouidian on-premises object storage, with scenario parameters summarized in Table 1.

Scenario Parameters	Value
Data storage capacity	1,000TB
% of data accessed per month	10%
Amount of data accessed	100TB
Analysis time horizon	5 Years

**Table 1. Scenario Assumptions for On-Prem vs. Public Cloud Storage Analysis**

Using the requirements and assumptions in Table 1, we configured a public cloud solution and a Clouidian on-premises solution and identified the costs associated with each one.

First, let's look at the public cloud option. The costs for this solution are listed in Table 2.

Expenditure	Public Cloud
Storage at rest (cost/month)	\$21,550
Storage access (cost/month)	\$7,800
Storage transactions (cost/month)	\$2,500
WAN bandwidth (cost/month)	\$5,000
Total Costs per month	\$36,850
<b>Total Cost for 5 Years</b>	<b>\$2,211,000</b>

**Table 2. Public Cloud Costs**

There are four separate fees that comprise the cost of public cloud storage. First, the user is charged a fee to store their data in the cloud (storage at rest). This fee varies based on the total amount of storage needed and will increase (or decrease) as storage volume increases (or decreases). Next, the user is charged a fee whenever they access their data (storage access). The more they access or manipulate their data, the higher the fee. A third charge is for the associated network bandwidth used (1Gb link) for this data access over the wide area. Lastly, the user is charged for certain storage operations performed (i.e. storage transactions — GET, PUT, LIST and more). This fee also varies — it depends on the number of operations performed in a given month.

Next, let's look at the costs of an on-prem object storage solution from Clouidian. These costs are listed in Table 3.

Expenditure	Clouidian On-Prem Object Storage
Hardware (with 5-year support contract)	\$450,000
Software (with 5-year support contract)	Included with hardware
Storage administrator (0.5 FTE, 5 years)	\$250,000
Power & Cooling	\$60,000
Rack space	\$15,000
<b>Total Cost for 5 Years</b>	<b>\$775,000</b>

**Table 3. Clouidian On-Prem Object Storage Costs**

The on-premises solution from Clouidian consists of four cost components. First, there is a cost for the hardware appliances and the associated software needed. For 1,000 terabytes, the hardware appliances require approximately 25% of a standard rack. Rack space costs are assumed at \$1,000 per month for a full rack, or \$250 per month for the space required here. Second, a storage administrator is required to operate the environment. For this analysis, one-half of a fulltime resource (.5 FTE) is needed at a cost of \$50,000 per year. Lastly, there is a cost for power and cooling for the equipment in the data center.

Once the costs of each storage option have been identified, the last step is to see how the total costs stack up against each other.

# Evaluating the Total Cost of Ownership (TCO)

When evaluating storage alternatives, the total cost of ownership over a set time period must be considered in order to obtain a complete view. For this analysis, we have used a five-year period. Table 4 lists the costs of the public cloud option and the Cloudian on-premise solution for comparison.

Expenditure	Cloudian On-Prem Object Storage	Public Cloud
<b>Hardware</b> (with support contract)	\$450,000	N/A
<b>Software</b>	Included with hardware	N/A
<b>Storage administrator</b>	\$250,000	N/A
<b>Power &amp; Cooling</b>	\$60,000	N/A
<b>Rack space</b>	\$15,000	N/A
<b>Storage at rest</b> (cost/month)	N/A	\$21,550
<b>Storage access</b> (cost/month)	N/A	\$7,800
<b>Storage transactions</b> (cost/month)	N/A	\$2,500
<b>WAN bandwidth</b> (cost/month)	N/A	\$5,000
<b>Total Cost for 5 Years</b>	<b>\$775,000</b>	<b>\$2,211,000</b>
<b>65% SAVINGS WITH CLOUDIAN SOLUTION</b>	<b>\$1,436,000</b>	

Table 4. TCO – Cloudian On-Prem vs. Public Cloud

When comparing costs associated with each alternative, Cloudian object storage offers an impressive 65% savings over the public cloud alternative, adding up to more than \$1.4 million over the five-year period.

When viewed side-by-side, it's easy to see the differences. While there is some upfront cost for hardware with the Cloudian solution, the monthly expense of storing, accessing, and manipulating data in the public cloud far outweighs the cost of equipment over the time period. In addition, it's important to note that egress fees (the cost to access or move data) in the public cloud, as well as associate network bandwidth costs, can be substantially higher than the amounts listed—if the user needs to frequently access a higher percentage of their data. In contrast, there are never any egress and bandwidth fees with the Cloudian on-premise solution. Users can access as much data as they want as often as they want without incurring fees.

While cost savings is a significant advantage of the Cloudian solution, it's not the only one. There are many more benefits to on-premise object storage.



# Additional Benefits of On-Premises Object Storage



As mentioned at the beginning of this paper, data fuels our world. From research and discovery, to workplace innovation and marketplace differentiation, to lifestyle conveniences, we rely on data and technology to live and to work. A byproduct of this reality is that data volumes are growing—rapidly.

Managing this growth in the most effective way possible is the quest of every organization. Fortunately, on-premises object storage has distinguished itself as a reliable, high-performing technology and offers many benefits that can make the job of managing data easier.

## Limitless Capacity

Object storage is scalable, providing limitless capacity. Unlike traditional disk-based storage systems, object storage is not hierarchical in structure. Nodes can be flexibly and non-disruptively added whenever additional capacity is needed, thereby eliminating size limitations and removing barriers to growth.

## High Data Durability

Protecting enterprise data against the risk of loss is important to every organization. Object storage provides high data durability by replicating or distributing data across multiple nodes or locations. Using either replication or erasure coding, storage can be flexibly configured to provide assurance against the loss of a single node or even a complete site.

## Consistent Security

Data breaches can be devastating to an organization. The need for strong data security, therefore, is obvious. Unlike the public cloud—where a configuration mistake can result in data being exposed to the public—on-prem storage is behind the firewall, and governed with a single security framework.

## Fast Data Recovery

Trying to recover data from the cloud can be painfully slow. If the data files being recovered are too large or too numerous, it may be nearly impossible to meet recovery time objectives. On-premises object storage eliminates this problem. It eliminates the impact of WAN latency, thereby enabling fast data recovery, even for large data sets.

## High-Performance

On-premises object storage provides faster throughput than the public cloud. For example, transferring a 1 terabyte of data over the wide area using a 1Gb link requires three hours to complete. In the data center, where there may be multiple high-speed links operating at low latency, the same operation would take minutes.

# Conclusion



In today's data-intensive world, organizations face many challenges. One such challenge is how to handle the growing volume of unstructured data. Should it be kept on-premises or moved to the public cloud?

Fortunately, the analysis is clear. When compared side-by-side, on-premises object storage offers the same or better capabilities than the public cloud in terms of functionality. And it does so for a fraction of the cost.

After careful review of the capabilities and analysis of the total cost of ownership of each option, on-premises object storage is more than just a

compelling alternative to public cloud. It is the most cost-effective solution for enterprise storage capacity.

It's easy to understand why object storage is quickly becoming the standard for capacity-intensive storage, backup, archive and recovery. And even with costs aside, on-premises object storage provides enterprises with the flexibility they need to respond to storage requirements, data security and control, and more.

If you're using public cloud for data storage, it makes a lot of sense to evaluate the benefits of switching to Clouidian object storage. With Clouidian, it's possible to benefit from both favorable economics as well as performance, security and control.

## Try Clouidian At No Cost Before You Start

Download the full-featured free trial of Clouidian HyperStore software and install it on any commodity hardware to build and test an on-premises object storage solution

[cloudian.com/free-trial](https://cloudian.com/free-trial)

### References:

<sup>1</sup> Patrizio, Andy. "IDC: Expect 175 zettabytes of data worldwide by 2025." Article on [www.networkworld.com](http://www.networkworld.com), December 3, 2018. Accessed March 14, 2019: <https://www.networkworld.com/article/3325397/idc-expect-175-zettabytes-of-data-worldwide-by-2025.html>

<sup>2</sup> Rizkallah, Juliette. "The Big (Unstructured) Data Problem." Article on [www.forbes.com](http://www.forbes.com), June 5, 2017. Accessed March 19, 2019: <https://www.forbes.com/sites/forbestechcouncil/2017/06/05/the-big-unstructured-data-problem/#e1da389493a3>

<sup>3</sup> Rowe, Cassandra. "The Trouble with Cloud 'Repatriation,'" 451 Research, August 3, 2018: Accessed March 11, 2019: <https://451research.com/blog/1960-the-trouble-with-cloud-repatriation>



**Clouidian, Inc.** 177 Bovet Road, Suite 450, San Mateo, CA 94402  
Tel: 1.650.227.2380; Email: [info@cloudian.com](mailto:info@cloudian.com); [www.cloudian.com](http://www.cloudian.com)