

2022-23 **DCIG** TOP 5 ON-PREMISES SDS OBJECT STORAGE SOLUTIONS



By DCIG Analyst, Todd Dorsey

On-Premises SDS Object Storage Solutions

Table of Contents

3	Object Storage Background
3	Object Storage Characteristics and Benefits
3	Public Cloud Object Storage Challenges
4	Why Object Storage On-Premises
5	Software-Defined Object Storage
5	Distinguishing Features of DCIG TOP 5 On-Premises SDS Object Storage Solutions
6	MinIO Inc. MinIO

On-Premises SDS Object Storage Solutions



SOLUTION

MinIO

COMPANY

MinIO Inc.
530 University Ave Ste B
Palo Alto, CA 94301
min.io

DISTINGUISHING FEATURES OF MinIO

- Deployment flexibility
- Performance focused
- Simplicity

DISTINGUISHING FEATURES OF TOP 5 SOLUTIONS

- Robust support
- Public cloud support
- Wide deployment support
- Directory services integration
- Robust analytics
- Value-added services

SOLUTION FEATURES EVALUATED

- Deployment capabilities
- Data protection
- Product and performance management
- Technical support
- Licensing and pricing

Object Storage Background

Object Storage has come a long way since its beginnings in the 1990s. Evolving from its modest role serving backup, archive, and disaster recovery needs, object storage has grown to become an essential part of enterprise storage. Object storage deployment options now abound: in the cloud, on-premises, and at the edge; appliance-based and software-defined; proprietary offerings and open-source solutions.

Object Storage Characteristics and Benefits

A principal reason driving object storage's adoption is the deluge of data growth. IT departments must store, manage, and protect existing data, plus deal with new data coming into the organization. The largest segment of this data growth is unstructured data such as video, images, documents, emails, presentations, spreadsheets, and similar file types. For many organizations, terabytes of data under management have become petabytes and even exabytes.

Object storage's characteristics are beneficial in this regard.

Scalability. There is a reason why all the public cloud providers primarily use object storage. Object storage's flat architecture without folders, hierarchical structures, or built-in limits does away with the scaling limitations of file systems. Plus, object storage scales at a more affordable price per terabyte than contrasting file storage systems.

Metadata. When storing large amounts of unstructured data, there is the need to make sense of that data. A key characteristic of object storage is the ability to attach customizable metadata to objects. Meta comes from the Greek word meaning 'with' or 'beside.' Metadata is a small amount of structured data associated with the object itself. User-defined metadata enables applications to perform bulk functions upon objects such as searching, indexing, and analyzing objects.

Self-protecting. Inherent within object storage solutions are protection capabilities to ensure data resiliency and durability. Often as a standard feature, object storage solutions include replication and erasure coding. Erasure coding segments and distributes data across multiple nodes or sites. Replication makes and stores multiple copies of the data for backup and recovery. Data may be restored to an original state from erasure-coded segments in the event of storage component failures, preventing data loss.

S3 API. Amazon's API for Amazon Simple Storage Service (S3), has become the API standard for object storage. All object storage vendors DCIG evaluated support the S3 API to some extent. With S3, organizations can access their object stores, authenticate data, obtain object properties, and manage permissions. With S3-compatible storage, organizations can move their data to other object storage services for cost savings or specialized workloads.

Public Cloud Object Storage Challenges

There is little question AWS's introduction of S3 cloud storage in 2006¹ fueled the adoption of cloud computing and storage. One of the great benefits of the cloud is its simple value proposition: with a credit card, scalable compute and storage infrastructure becomes immediately available. However, public cloud infrastructure is not without its challenges.

On-Premises SDS Object Storage Solutions

Cloud costs are a principal reason why organizations are repatriating their data to on-premises or hybrid-cloud (on-premises and cloud) object storage solutions.

Latency. Latencies occur because of transmission delays across the Internet or within the cloud provider itself. Latencies can be problematic for workloads with time-sensitive requirements. Dedicated circuits into the cloud provider can help, but may be insufficient for critical applications.

Data sovereignty. The rise of cloud computing has resulted in countries passing regulations that specify how data is stored. Regulations may require data collected and processed within a country's borders remain within those borders. It may also mean that data is subject to the jurisdiction in which the cloud infrastructure physically resides. Ultimately, enterprises must understand the evolving legal impacts of data sovereignty regulations on data storage and transfers.

Public cloud data breaches. The accessibility and volume of data stored in the cloud makes it a attractive target for criminals. A recent security research study reported that 40% of organizations had experienced a cloud-based data breach over the last 12 months.² Given the costs and negative public relations that data breaches can bring, organizations are rightly concerned about this potential risk.

Cloud cost overruns. While storing data in the public cloud is relatively inexpensive, frequently accessing that data adds up. Further, it is not always clear how cloud decisions will affect an invoice. It is a common experience for companies using the cloud to receive a high bill at some point. Cloud survey research showed a large percentage of respondents had experienced unexpected or significantly higher cloud costs due to COVID-19.³

Why Object Storage On-Premises

Certainly, there are good reasons enterprises should leverage the cloud for its convenience and strengths. The public cloud helps companies learn cloud skills and technologies. The 'however' is that the cloud does not scale cost-effectively. Rising cloud costs place negative pressures on margins. As the team at Andreesen Horowitz noted in their article *The Cost of Cloud, a Trillion Dollar Paradox*, 'You're crazy if you don't start in the cloud; you're crazy if you stay on it.'⁴ Cloud costs are a principal reason why organizations are repatriating the data to on-premises or hybrid-cloud (on-premises and cloud) object storage solutions.

Other reasons driving the adoption of on-premises object storage include:

Unstructured data growth solution. Object storage offers enterprises an optimum solution for dealing with exponential unstructured data growth. As organizations create unstructured data, they can move cold data off primary storage to on-site archive or active archive object storage. Enterprise may also use object storage for their backup and disaster recovery needs. Moving cold data off tier 1 storage creates a lean, organized, and higher performing primary storage infrastructure.

High performance object storage. Data rich applications and workloads depend on fast response times. Hardware and software improvements add to on-premises high performance use-cases involving unstructured data such as IoT, big data analytics, and AI/ML. Flash media accelerates low latency and high throughput. On top of this, because of proximity, workloads utilizing on-premises object storage achieve better performance.

On-Premises SDS Object Storage Solutions

Object storage becomes the preferred storage medium where IoT devices, cameras, and machine sensors generate, store, and transmit unstructured data.

Data sovereignty. Data sovereignty rules change yearly. Non-compliance can lead to stiff penalties. For critical data and workloads, organizations may find it simpler to store their object data on-premises for geographical control. On-premises object storage also allows organizations to design security protocols tailored to their needs.

Edge computing. Edge computing use cases continue to evolve. The edge presents unique challenges of environment, power, performance, connectivity, and space. Object storage becomes the preferred storage medium where IoT devices, cameras, and machine sensors generate, store, and transmit unstructured data.

Transition to cloud-native workloads. Enterprises are using on-premises object storage to develop, test, and deploy applications made for a cloud computing architecture. Cloud-native workloads run anywhere: in the cloud, in data centers, or at the edge. Developing these applications on-premises enables developers to work out issues before transitioning these workloads to the cloud. This on-premises storage also enables organizations to move applications developed in the cloud back on-premises.

Software-Defined Object Storage

The growth of Software-defined Storage (SDS) is part of the broader transition to the software-defined data center, where infrastructure elements such as compute, storage, and networking are abstracted and virtualized. Increasing SDS capabilities, along with its flexibility and cost-efficiencies, fuels SDS demand.

SDS expands on the benefits of object storage in the following ways:

Scalability. SDS solutions enable organizations to scale up or out depending on requirements. On-premises, organizations can add to their existing physical storage to expand their virtualized pool. Many SDS solutions have virtually no capacity limitations. Such scalability helps organizations flex to their growing data and application needs.

Cost efficiencies. SDS-based solutions bring savings and efficiencies compared with traditional storage systems. Organizations may deploy SDS solutions with cost-effective commodity hardware. They may also save money by optimizing and extending the life of their existing storage. When constrained IT budgets rarely grow at the pace of unstructured data accumulation, these cost-efficiencies are greatly valued.

Data management capabilities. The data management component of SDS solutions virtualizes and views data from its source systems as one central repository. Global views, including permissions management, capacity utilization, and analytics, enable new opportunities for ensuring optimal performance and cost for managing an organization's object stores.

Distinguishing Features of DCIG TOP 5 On-Premises SDS Object Storage Solutions

DCIG evaluated eighteen SDS-based solutions for an on-premises object storage use case. Using feature-based analysis and comparisons of defensible data derived from publicly available sources, vendors, and DCIG's own experience, DCIG's TOP 5 On-Premises SDS Object Storage Solutions evidenced these characteristics in contrast with the other evaluated solutions.

On-Premises SDS Object Storage Solutions

DCIG TOP 5 solutions offer ample deployment options compared with the other evaluated solutions.

Robust support. DCIG TOP 5 providers display robust support capabilities. All TOP 5 vendors provide 24x7x365 technical support and one-hour support response times. Each of the TOP 5 solutions supports real-time reporting to the solution vendor for expedited trouble reporting and resolution compared with 69% of the other evaluated solutions. All vendors offer community support forums and knowledgebases for self-service support. In short, DCIG TOP 5 solutions evidence a greater breadth of support options than the other solutions.

Public cloud support. Enterprises may leverage public cloud storage for its benefits. This may be for backup, disaster recovery, or archive purposes. DCIG TOP 5 solutions support multiple public cloud providers. Such broad support offers flexibility in matching a cloud provider's cost and capabilities with the needs of the business.

Wide deployment support. DCIG TOP 5 solutions offer ample deployment options compared with the other evaluated solutions. Examples include deployment support on bare metal servers, commodity x86 servers, Linux operating systems, public cloud, and VM environments. Wide deployment options give organizations flexibility for using SDS Object Storage for their specific situations.

Directory services integration. DCIG TOP 5 solutions show broad support for directory service and authentication features. All support AD/LDAP integration, IAM authentication, and object level access control lists for data security. In contrast, only 34% of the other evaluated solutions provide IAM and Object-Level Access Control List support.

Robust analytics. DCIG TOP 5 Solutions deliver extensive analytic capabilities. Analytics give organizations insights into their data patterns and trends. These insights help IT decision-makers plan, budget, and optimize their object storage usage.

Value-added services. In addition to their SDS Object Storage offering, DCIG TOP 5 Solution providers also offer value-added services such as consulting and installation services.

MinIO Inc. MinIO

MinIO's founders created MinIO to do one thing very well: serve objects fast.

Upon DCIG's completion of reviewing multiple, available SDS-based on-premises object storage solutions, DCIG ranked MinIO as a TOP 5 solution. MinIO's founders created MinIO to do one thing very well: serve objects fast. Because MinIO is purpose-built to serve only objects, a single-layer architecture achieves all necessary functionality. MinIO's open source solution deploys across a wide variety of platforms and hardware. Enterprises may manage their MinIO object storage through the MinIO Client (command line), MinIO Console (GUI), or MinIO Operator (Kubernetes). MinIO provides straightforward licensing with a no-charge community support plan, a standard plan, and an enterprise plan. Priced on capacity and billed monthly, one license covers the storage suite. Paid licenses include SUBNET, MinIO's subscription network offering of 24x7x365 engineering support, architecture and security reviews, and health diagnostic capabilities.

Notable features that earn MinIO a DCIG TOP 5 award include:

MinIO everywhere. Kubernetes-friendly and cloud-native, MinIO supports diverse deployments from edge locations to the cloud. MinIO is S3-compatible, but can be used in all non-S3 public cloud platforms such as Azure, Google Cloud Platform, or Alibaba and has native marketplace offerings for AWS, GCP and Azure. Additional features including multitenancy, replication, IAM controls, and at-rest and in-flight encryption make MinIO ideal for on-premises, hybrid cloud, and multi-cloud solutions.

On-Premises SDS Object Storage Solutions

Performance focused. MinIO regularly publishes performance and comparison testing on the MinIO website. Uniquely, MinIO writes object and metadata together to improve performance and especially speeds workloads involving small objects. By employing Single Instruction Multiple Data (SIMD) acceleration, enterprises experience fast performance even on commodity hardware. MinIO's focus on serving objects fast especially benefits modern use cases requiring high throughput such as AI/ML and analytics.

Simplicity scales. MinIO argues that simple scales, both technically and operationally. For example, MinIO's small binary (~100MB) installs within minutes by downloading a Docker image or a single binary and then executing it. Its lightweight footprint allows it to run on small edge devices, scale to multi-node storage architectures, and power multicloud storage solutions. MinIO's commands, configurations, and variations are kept to a minimum, reducing administrative tasks and the risk of errors. ■

Sources - Referenced in January 2022

1. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/s3-userguide.pdf>
2. <https://www.securitymagazine.com/articles/96412-40-of-organizations-have-suffered-a-cloud-based-data-breach>
3. <https://www.informationweek.com/cloud/10-cloud-strategies-to-avoid-cost-overruns>
4. <https://a16z.com/2021/05/27/cost-of-cloud-paradox-market-cap-cloud-lifecycle-scale-growth-repatriation-optimization/>

About DCIG

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analysis. DCIG analysts provide informed third-party analysis of various cloud, data protection, and data storage technologies. DCIG independently develops licensed content in the form of TOP 5 Reports and Solution Profiles. Please visit www.dcig.com.



DCIG, LLC // 7511 MADISON STREET // OMAHA NE 68127 // 844.324.4552

dcig.com

© 2022 DCIG, LLC. All rights reserved. Other trademarks appearing in this document are the property of their respective owners. This DCIG report is a product of DCIG, LLC. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. Product information was compiled from both publicly-available and vendor-provided resources. While DCIG has attempted to verify that product information is correct and complete, feature support can change and is subject to interpretation. All features represent the opinion of DCIG. DCIG cannot be held responsible for any errors that may appear.

Licensed to MinIO Inc. with unlimited, unrestricted, global distribution rights through December 31, 2023.

February 2022 7