

White Paper

Cloudian HyperStore: Customer Insights and Best Practices

Sponsored by: Cloudian

Amita Potnis
August 2018

EXECUTIVE SUMMARY

Many organizations are at a turning point as they strategize their storage infrastructure needs for the long haul. In many interviews with IDC, end users reveal similar concerns consistently. Unsurmountable data growth, risks, long-term data retention for compliance, and cost constraints are some major concerns that define the road map for IT infrastructure within any organization. Generally, a road map is a set of requirements that drive the adoption of certain technologies or platforms. For example, rightsizing applications may mean potentially adopting newer platforms (private/public cloud, IaaS, PaaS, and on-/off-premises traditional storage). IDC's 2018 *Data Services for Hybrid Cloud Survey* indicates that while organizations are adopting a cloud-first strategy, security concerns (primarily for data in use) lead them to deploy a certain percentage of the workloads on-premises, driving the overall hybrid cloud adoption. The same survey indicates that data loss prevention, ensuring data quality, and regulatory compliance are the top data-related challenges overall.

Specifically with regard to unstructured data, IDC sees an increased shift to adoption of object-based storage (OBS). According to *File- and Object-Based Storage Survey Findings, 2017: Adoption and Workloads Trends – Part 1* (IDC #US43630018, March 2018), increased management complexity and high maintenance costs are leading factors to either replace or augment network-attached storage (NAS) with OBS. In addition, as organizations adopt a hybrid cloud strategy, data management and tiering to various platforms (on-/off-premises traditional storage or cloud) through a single pane of glass are becoming necessities rather than nice-to-have features. Native file system capabilities, comprehensive information life-cycle management (ILM) policy capabilities, and a solution-centric ecosystem of technology partners are top of mind when selecting an OBS vendor.

Vendors and end users alike should bear in mind that customer experience (CX) is becoming increasingly important. It is clear that having a "best in class" feature set is no longer enough to win business. Ease of procurement, simplistic billing, nondisruptive technology refresh, predictive analytics and maintenance, availability guarantee, self-healing, support for automation, and predefined workflows for ease of use are some of the CX defining criteria.

As OBS solutions become mainstream and begin supporting production workloads, all the concerns mentioned previously become critical to vendor/platform/solution selection. IDC interviewed three customers that shared their experiences and insights regarding their choice of Cloudian as an OBS provider. IDC believes that Cloudian has come a long way in terms of making its mark and proving its longevity as well as demonstrating a steadfast focus on product development. Cloudian's customers reiterate all of the abovementioned concerns and give a detailed account of the current deployment and the eventual outcome. IDC finds that Cloudian's customers overall are not just pleased with the feature set supported by HyperStore but also satisfied from a CX perspective.

CLOUDIAN'S OBS OFFERING

Cloudian has been a long-standing object storage supplier offering an object storage platform called Cloudian HyperStore and an enterprise file services offering called Cloudian HyperFile. Through these two offerings, Cloudian aims to consolidate file and object data into a single repository. Cloudian HyperStore is an on-premises, scale-out object storage system with a multitenant, multiapplication, and distributed storage technology that can scale from 10TB to large-scale environments comprising hundreds of petabytes.

Typically deployed on-premises by organizations building a private cloud, HyperStore is also widely used by public cloud service providers to support S3-compatible storage services. The product is available in the form of an appliance or as software that can be deployed on commodity hardware. Cloudian's high-density appliances support up to 168TB in a 1U-high device or 980TB in a 4U-high device. Additional appliances, or "nodes," can be added seamlessly to a storage cluster without downtime, and appliances can be refreshed without disruption. This combination allows business to run without any interruptions. Cloudian HyperStore provides a single pool of storage for reduced operational expense, enabling a full-time IT employee to manage installations larger than 10PB single-handedly.

In addition to full S3 API compatibility and petabyte scalability, a key differentiator for Cloudian HyperStore is the company's hybrid cloud data management capability — it can tier data to the public cloud by either replicating or migrating on-premises data. In either case, Cloudian manages the entire data set, on-premises and in the cloud, from a single pane of glass. Multicloud support is included as well, including Google Cloud Platform, Microsoft Azure, Amazon S3, and S3-compatible public clouds such as NTT's Cloud-N or Interoute's GDPR-compliant cloud. In each of these clouds, users can read/write data using the same S3 API. In each cloud, data is stored in the platform's native format, allowing, for example, data stored to Azure Blob to be directly accessed by Azure applications, without the need to traverse the HyperStore platform.

With Cloudian, organizations therefore have a single S3-compatible access point and a single management console that can traverse all major cloud platforms and that will store data in that cloud's native format. More recently, Cloudian expanded its data management capabilities and support for new S3 APIs such as S3 Select. Cloudian offers granular, bucket-level, policy-based asynchronous and synchronous replication and erasure coding, alongside other features such as strong and eventual consistency. Multitenancy is built with features such as quality-of-service (QoS) controls, role-based access control (RBAC), billing, and reporting.

Cloudian acquired Infinity Storage, an Italian file-based storage supplier, and has integrated the offering with HyperStore in the form of HyperFile. HyperFile supports enterprise NAS features such as snapshots; write once, read many (WORM); nondisruptive failover with active-active controllers; and bimodal data access (write with file, read via S3, and vice versa). It supports nondisruptive data migration from existing NAS, with policy-based migration and no service interruption from start to finish. HyperFile supports a single namespace across sites, configurable consistency, and client-level security and authentication. Both HyperStore and HyperFile are available as software-only solutions as well as turnkey integrated appliances. Cloudian's HyperStore employs a peer-to-peer architecture in which all nodes are active and can perform all functions (e.g., access and storage), and therefore, most functions, if not all, are performed in parallel. Read, write, replication, repair, and so forth of large objects that are sliced across multiple nodes are performed in parallel across multiple nodes.

Cloudian claims that over the past 12 months, it has doubled its customer base to over 200 buyers, increased its repeat customer revenue by five times, inked a worldwide reseller deal with HPE, and established a presence in France, Germany, Italy, the Nordics, and Benelux. The company also announced a \$125 million joint venture with Cisco-backed private equity firm Digital Alpha to accelerate Cloudian's market penetration via a Cisco partnership and a consumption-based procurement model.

CUSTOMER INTERVIEWS AND INSIGHTS

IDC conducted in-depth interviews with three customers of various sizes to understand how Cloudian is supporting their specific storage needs today and their expected expansions of current deployments and potential outcomes. The customers will remain anonymous in this study but will provide an objective interpretation of real-world Cloudian deployments, best practices, and advice. Insights from the interviews are discussed in the sections that follow.

Customer 1 – Large Retail Chain Store in the United States

Overview

IDC interviewed a large retail chain store based in the United States. This organization has over 100,000 employees with over 100 IT staff members that support platforms, storage, networking, virtual machines, and end-user compute (desktops/virtual desktop infrastructure, voice/telecom, etc.). Overall, the organization runs over 6,000 applications with over 2PB of unstructured data housed on NAS arrays. The organization has specific data retention policies spanning 17 days to 20 years based on data type. For example, structured data is retained for 17 days and unstructured data is held for 35 days; for tax-compliance purposes, data can be held for 90 days to 20 years. IDC would like to highlight that this organization has owned legacy object stores.

Requirements

The use of smartphones and increase in online shopping made significant changes to how business operated for this organization. Products sold in this organization's local stores needed to be listed in high-resolution images on the organization's website to facilitate revenue streams from online shopping. One of the biggest challenges resulting from this change in business was the significant increase in high-resolution media files (product photos and videos). At the time, these high-resolution media files were housed on expensive NAS, causing cost and performance overheads. To organize, adapt, and distribute this rich media, the organization uses S3-compliant Adobe Experience Manager Assets (EMA) as its digital asset management solution.

The organization believed that its requirements of a new storage solution were beyond just scale and information life-cycle management. The company needed the flexibility of a platform that could support traditional and next-generation workloads. The new solution needed to be backward and forward compatible as the organization tried to retire/replace its legacy object stores and backup appliances. The company also needed the new product to be simple to use (as opposed to excessive command-line scripts), with the ability to offer its users a self-service model. Security was another key concern. One of the key criteria for any product to be considered was that it should have the ability to support security certificates, multitenancy (support for LDAP), encryption, and so forth. In addition, the organization had specific goals to adopt public cloud storage services from providers such as Microsoft, Amazon, and Google. Given the corporate strategy to adopt public cloud storage, the new storage solution needed to have the ability to run and support Adobe EMA on-premises or in the cloud, thus maintaining consistency in user experience. Finally, the new storage solution had to be cost effective.

The organization evaluated three object-based storage solutions and, after assessing business and technical requirements and conducting a proof of concept (POC), decided to deploy Clouidian's HyperStore object storage solution.

Solution and Future Plans

More than 18 months ago, the organization deployed Clouidian's minimum configuration of 100TB across two sites. Part of the reason for opting for the minimum configuration was that the senior leadership within this organization was new to OBS and wanted to ensure that this was the right choice before expanding the footprint. After the initial purchase of a minimum base system, the organization expanded its OBS footprint to 450TB in a short span. The organization procured Clouidian's appliance because it wanted to have a single point of contact for any component failure – hardware or software. Currently, Clouidian houses the organization's high-resolution print and online content managed by Adobe's EMA.

The organization has plans to replace existing storage backup arrays with Clouidian to support long-term retention and backup. In addition, the organization aims to deploy Clouidian behind its existing all-flash array to serve as an active archive repository for unstructured data. For 2018, the Clouidian deployment is expected to grow to 1PB. The organization is also considering Clouidian's HyperFile along with HyperStore as a unified file and object repository as opposed to procuring and maintaining another purpose-built appliance.

Outcome

Deploying Clouidian's HyperStore was advantageous to this large retail chain store in many ways. The organization reported the following outcomes of deploying Clouidian:

- **Support for traditional and next-generation use cases:** Clouidian supports backward and forward compatibility. For example, support for CAS allows easy migration of data housed in legacy object stores such as Atmos and Centera.
- **Cost savings:** The organization believes that Clouidian has offered better affordability, with close to half a cent per gigabyte per month for a single site and one cent per gigabyte per month for replication to two sites for an active-active deployment needed for Adobe EMA. The organization expects prices to keep going down as hardware prices decrease.
- **Management buy-in:** Senior IT leadership now understands how OBS works and which use cases it supports and is convinced of its benefits and advantages, which has led to the expansion of Clouidian's footprint in 2018.
- **Technology partnerships and ecosystem:** Clouidian's HyperStore is certified to run Rubrik and Nutanix.
- **Experience:** The organization claims that it has experienced good technical and sales support, with few issues or enhancements that had to be addressed, along with easy procurement and rolling upgrades with no service disruption.

Customer 2 – Large Broadcasting Studio in the United States

Overview

IDC interviewed a large broadcasting studio based in the United States with over 50,000 employees and a run rate of over \$25 billion. The organization houses over 6PB of storage including block, file, and object data and is growing at about 500TB a year. The organization heavily uses LTO tape libraries to archive its digital assets but had goals to deploy new storage technology that would serve as an active archive tier to keep data accessible to its users at all times.

Requirement

Digital transformation changed business process for this broadcasting studio, much like it did for the large retail chain store. Data, or "media assets" in the media and entertainment world, is no longer played on just the television or radio. End users choose to consume media on various smart devices. Video assets need to be transcoded to various formats to run on these various devices, and much of the consumption is on demand; therefore, assets need to be accessed and retrieved easily. This was not possible with the organization's use of LTO tape libraries. In addition, the affordability of tape libraries was quickly diminished by other overheads. While tape media afforded the organization a significantly lower dollar-per-gigabyte ratio compared with its high-performance NAS arrays, one of the things that was a constant challenge was the cycle of migrating assets from one older generation of LTO to the latest. This process was prolonged and multifold. IT personnel must first manually allocate hard disk drive (HDD) space on external storage infrastructure, then move data from legacy LTO to HDD and, finally, move the data on HDD to the latest generation of LTO tape library. This meant a painful expenditure of engineering time at all levels – process, hardware, and software, along with service windows for operation tasks.

This led to the organization investigating viable OBS solutions that would meet the key requirements of low cost, accessibility, and scalability with minimal overhead. The organization conducted a 10-year total cost of ownership (TCO) exercise for OBS and compared the results with public cloud storage services, which have a lower capital expenditure model. The company's findings indicated that 5-10 years of public cloud storage services usage was unsustainably expensive for the company's needs, while most other commercial OBS providers had high year one out-of-budget expenditure. Cloudian was the only solution that fit the needs of this large broadcasting studio house in terms of not just cost savings but also scalability, low maintenance, business continuity, and so on.

Solution

The organization deployed an 18-node Cloudian HyperStore appliance that was slated to replace its entire LTO tape library. The organization chose to maintain and deploy a separate NAS solution for its editing and transcoding needs. The solution serves as an active archive tier with 4PB of the organization's 6PB of total storage capacity on Cloudian across two sites. A second copy of the data assets is replicated to the second site. Currently, the solution supports about 250 internal users. The organization is not using Cloudian's HyperFile offering today but may consider it in the future for non-mission-critical workloads.

Outcome

- **Experience:** The organization took 14 days to deploy the Cloudian cluster from the time product was received, mainly because it had to be racked and wired. The company claims that moving digital media assets to Cloudian HyperStore from LTO tape libraries significantly reduced time and head count requirements for maintenance and support. For example, the organization does not have to invest manpower in swapping out LTO and bringing technicians onsite to repair failed mechanical components.

- **Cost savings:** Initial expenditure for Cloudian including support cost for 4PB in a single datacenter was less than 50% of the cost of an offering from a competitive vendor and AWS. In addition, the TCO for a 10-year period was much lower than the TCO for storage services offered by Amazon.
- **Low maintenance and overheads:** The organization claims that the only maintenance that it has to manually cover is component (hardware) failure owing to over a 100 spinning disk drives that are currently deployed in the Cloudian cluster.
- **Business continuity:** Cloudian's automatic data verification and self-healing system provides resilience and reliability. The system is highly redundant and therefore avoids downtime when replacing or upgrading systems. The organization has experienced 0% downtime and no production impact owing to rolling upgrades.
- **Centralization/standardization:** Standardized Asset Management Technology Stack implemented across OBS environments allows archived data repositories to be federated rather than siloed.

Customer 3 – Wealth and Asset Management Services Company in Europe Overview

IDC interviewed a wealth and asset management services company in Europe that employs more than 4,000 employees and hundreds of investment managers. The company was looking for a consolidated file and object storage solution for its unstructured data storage needs. Being a wealth and asset management company, the organization has access to and owns sensitive data that cannot be pushed into the public cloud because of compliance reasons. In addition, the organization needed compatibility with storage in the cloud for nonsensitive data. Currently, the organization houses close to 6PB of data across siloed storage solutions for block, file, and object.

Requirement

Unlike most customers that strategize to procure storage solutions based on existing needs, this organization presents a unique and different case. Expecting data growth and possible net-new greenfield deployments, this organization decided about a year ago to investigate a new storage platform that supported both file and object storage in a single solution. The idea was to streamline its storage needs with a unified file and object offering and do away with storage silos. At the time this expansion was planned, OBS solutions with file system capabilities were unable to do a POC because of limited clarity on which workloads the new solution will support or how the solution will support those workloads. The organization therefore decided to procure and deploy two separate solutions: a solution that would support high-performance file-based workloads and another OBS solution.

The organization was impressed with Cloudian enough to conduct a POC with a little clarity on use cases in a small environment using six virtual machines that deployed in less than a day. The organization evaluated several OBS solutions and found Cloudian's S3 compatibility to be superior than others. At the time of the company evaluated Cloudian, HyperFile capabilities were not generally available, and therefore, the organization chose to deploy separate solutions in support of file-based storage and OBS. The organization is pleased overall with the Cloudian deployment but thinks that load balancing should be integrated into the product and that it should also allow bucket-level policy changes.

Solution

In August 2017, the organization procured and deployed six Clodian HyperStore 1508 appliances across two sites amounting to a total of 864TB. Clodian serves as a production environment with deployments across two sites using replication over erasure coding for an added layer of data protection. The organization uses Clodian to support two production workloads: on-premises CloudStack to store status records as well as archive data (before being deleted from production environments) and a sandbox test and develop three-cluster environment.

For less sensitive data, the organization will investigate using Clodian for private cloud deployments with replication to the public cloud. As the organization continues to execute its hybrid cloud/multicloud strategy, it will look to cross-replicate data across on-premises Clodian HyperStore OBS and public cloud storage services such as Amazon or Google.

Outcome

- **Planning for the unknown:** As hard as it is to plan for known requirements, it is exponentially harder to plan for the unknown. The organization appreciates Clodian's ability to conduct a POC under the circumstances.
- **Ease of management:** The organization reports that the Clodian web interface is easy to use with dashboards that provide metrics on resource consumption by application, billing, and so forth.

RECOMMENDATIONS AND ADVICE FOR BUYERS

Buyers of any new storage technology must think about the following before making any investments in procuring infrastructure:

- Analyze how your business has changed in the past few years and how it is expected to change over the next two years. This will help define the type of infrastructure needs.
- Define your requirements as best as you can for today and tomorrow.
- Which features are must-haves and therefore deal breakers (e.g., encryption)?
- Do you need a solution with native file system capabilities on OBS?
- Do you need support for flash for improved performance?
- What license schemes are available that best suit the needs of your organization?
- What level of insights would you like to see in terms of application and resource utilization and management?
- Ask for POC and customer references for ease of use and best practices.
- Look for vendor commitment to technology development, maintenance, and support as well as go-to-market strategy and financial stability.

CONCLUSION

Clodian has made strides in recent months to add significant capability to its OBS solution. If the company continues to deliver on its promise in terms of not just development but also go-to-market strategies and extensive channel and technology partnership programs, the product is slated for success.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2018 IDC. Reproduction without written permission is completely forbidden.

