

7 Reasons to Run Data Analytics on a Cloudian S3 Data Lakehouse

The modernization of the data analytics architecture that started in the cloud combines the flexibility, cost-efficiency, and scale of data lakes with the data management and ACID transactions of data warehouses giving us the Data Lakehouse. Data lakehouse allows enterprises to analyze large amounts of structured, semi-structured, or unstructured data, for instance, web server logs, RDBMS data, NoSQL data, social media, sensors, IoT data, and third-party data, all stored in one place, in the same format as its source systems in a modern S3 data lake.

Here are 7 reasons why Cloudian HyperStore is the best option for enterprise customers looking to modernize their data analytics with a data lakehouse architecture:



1. It Brings the Cloud to Your Data Center

Modernization of data analytics started with the cloud. However, while most organizations want to be cloud-centric, only a fraction of the data can reside in the cloud. Cloudian HyperStore brings the scalability and flexibility of the cloud into your data center as part of a hybrid cloud data lakehouse solution built on S3, enabling modern data analytics on-premises. Separate compute from storage scalability and expand by simply adding nodes, in one site or across multiple sites. Manage it all within a single namespace, from a single management console, and search metadata across all your sites with a single query. Integrated tools even let you create a data copy to a public cloud for disaster recovery, if desired.



2. Easy Integration with All Leading Analytics Platforms

Cloudian HyperStore, the best-in-class on-prem hybrid cloud storage system with native support for S3 API is validated and certified with leading data analytics platforms, including Greenplum, Teradata, Vertica, Apache Druid, Microsoft SQL Server 2022, Splunk, Elastic, and Cribl, amongst others. Only Cloudian offers a 100% native S3 API object storage system to modernize data analytics architectures on-premises. Together with these industry-leading analytics applications, Cloudian provides customers with a robust data lakehouse solution that includes scalable storage, high data durability, and fast recovery times, within the security of your own firewall, at up to 70% lower total cost than traditional storage solutions.



3. The Data Lakehouse for All Use Cases

Data lakehouses built on Cloudian Hyperstore can be used as storage destinations for any application built for the S3 API. In addition to data analytics, organizations can use HyperStore as the repository for machine logs, media files, remote directory, as well as for a secure backup target (databases and/or system-level) with integration into Veeam, Rubrik, Commvault, Veritas, HYCU, and any other application vendors that support the S3 API. HyperStore brings the data to the application, for more efficient storage, management, and security.



4. The Most Security Certifications of Any On-Prem Object Storage

Cloudian HyperStore provides extensive security features that make it possible to deploy and operate a cost-effective data lakehouse that keeps organizational data extremely secure. Cloudian HyperStore has more security certifications than any other on-prem object storage. Our DOD-level security certifications even make it possible to operate a secure AWS Outposts deployment. Capabilities include AES-256 server-side encryption for data stored at rest, SSL for data in transit (HTTPS), transparent key management, Amazon Identity and Access Management (IAM), API and role-based access controls (RBAC), audit trail logging, secure shell with root disable, and last but certainly not least, S3 object lock to keep your data immutable against the threat of ransomware.



5. High Data Resiliency

Data lakehouses built on Cloudian Hyperstore provide 14 nines' of resiliency. HyperStore gives the option of storage policies (administrator selectable) for implementing data-protection based Replication (RF) or Erasure Coding (EC). Administrators can configure the number of replicas or type of erasure code scheme required to meet SLA and cost objectives. Storage policies also provide fine-grain control of data placement across data centers, taking into consideration factors such as cost-efficiency, security levels, and proximity.



6. Hybrid and Multi-Cloud Ready

Cloudian HyperStore's native S3 API cloud integration gives data lakehouses the power to tier data across on-prem storage and private/public clouds. This lets organizations cost-effectively combine storage across environments into a single pool while consolidating storage management enterprise-wide to a single screen. This flexibility is ideal for cloud bursting—a temporary use of more storage resources than usual—or disaster recovery (DR) in public cloud.



7. White Glove Support

Cloudian provides a unique combination of tools and expertise to help you get the most from your Cloudian deployment, whether single location or petabyte scale-out across multiple locations around the world. Cloudian's support offerings are designed to ensure the success and operational health of your Cloudian deployment, maximizing uptime and performance while also delivering proactive analytics and insight at up to 70% less cost than traditional storage options.

[Learn more at Cloudian.com/Data-Analytics.](https://cloudian.com/Data-Analytics)



Award-Winning

Proven at over 700 enterprise customers worldwide with nearly two exabytes of capacity under management, Cloudian Scored Highest for All Use Cases in the Gartner 2020 Critical Capabilities for Object Storage Report and was named a Gartner Peer Insights Customers' Choice in 2020, 2021, and 2022.