

# TOP 12 REASONS WHY CUSTOMERS CHOOSE CLOUDIAN HYPERSTORE®



## 100% S3

Simple Storage Service (S3), pioneered by Amazon, has become the “de facto” standard for object storage in the cloud. Companies and developers implementing S3 apps that may run on-premises, or in hybrid/private cloud should consider Cloudian HyperStore storage system as it is the only object-storage platform 100% fully-compatible with Amazon S3’s HTTP REST APIs. It even reports the same error codes as AWS. Customers’ existing HTTP S3 applications are guaranteed to work and can even use the same AWS S3 SDK for building S3 Apps.



## Multi-tenancy & Self Service

The Cloudian HyperStore storage system provides the capability to have multiple users reside on a single, shared infrastructure without compromising security. Data for each user is logically separated from other users’ data and cannot be accessed by any other user unless access permission is explicitly granted. Cloudian HyperStore gives role-based access to system and group administrators and to users. Users can select and provision storage services on-demand from a service catalog, all done via a self-service portal.



## QoS & Quotas

Cloudian HyperStore system administrators can set usage rate limits and storage quotas on a per-group and per-user basis. Also group administrators can set rate controls and quotas for individual members of the group thus allowing prioritization and ensuring preferred access to resources for select users within a group. This level of granular control also ensures that resource consumption (BW, Capacity) does not go overboard for certain users or groups beyond the limits of the Cloudian HyperStore storage system.



## Protect & Distribute Data with Storage Policies

Cloudian HyperStore provides various storage policies (administrator selectable) for implementing data protection based on value of the data. Users when they create a new storage bucket have the ability to choose which pre-configured storage policy to effect in order to protect the data. Hyperstore allows customers to protect and distribute data using Replication, Erasure Coding and Replication over Erasure Coding. With replication, a configurable number of copies of each data object are maintained in the system, and each copy is stored on a different node. Copies can be at various sites if desired. With erasure coding, each object is encoded into a configurable number (known as the “k” value) of data fragments plus a configurable number (the “m” value) of redundant parity fragments. Each fragment is stored on a different node, and the object can be decoded from any “k” number of fragments. Replication over EC can also be used to replicate data across DCs for Disaster Recovery (DR) while maintaining local erasure-code fragments. Using Replication over EC within a DC, eliminates unnecessary data movement over the WAN in the event of Node or Disk failures. Storage policies also enable customers to get fine grain control of data placement across data centers taking into consideration various factors such as cost efficiency, security levels or proximity.



## Scale Out with Low Cost Commodity Hardware

Running on off-the-shelf commodity hardware, a Cloudian HyperStore system can scale up to thousands of nodes across multiple datacenters, supporting millions of users and hundreds of petabytes of data. New nodes –heterogeneous if desired– can be added without service interruption. Performance, capacity and resiliency also increase in the process. Each node is an S3 server and can provide data access across the cluster, resulting in a true peer-to-peer platform with no single point of failure or bottleneck.



## Hybrid Cloud Tiering (Tier to Other S3 Platform)

With Cloudian HyperStore, you can define a tiering policy per bucket. Objects in that bucket can be tiered to Amazon S3 or Glacier, a different HyperStore cluster, or an S3 compatible off-line storage system such as tape. Warm Data can be stored on Cloudian, cold data can be stored on tape, etc.



## Smart Operations

The entire Cloudian HyperStore system, potentially spanning multiple regions, sites and nodes, can be managed easily from a single CMC (Web) console. The Cloudian HyperStore has the ability to automatically perform system tasks to ensure that each replicated data object in the cluster has the proper number of replicas, that each copy is up-to-date and that each copy passes for integrity check. If erasure coding (EC) is in used, it's also important that erasure coded data be regularly evaluated and repaired to ensure that for each object, the right number of fragments exist across the cluster and that all fragments are up-to-date. Also, built-in to the Cloudian HyperStore storage system is Smart Support. It is turned on by default.

Every deployed Cloudian HyperStore storage system feeds telemetric data to Cloudian Analytics platform (Hadoop) that mines the data and pre-emptively alerts customers for pending or potential issues. A case number is automatically generated by the system, and Cloudian Support works with the customer to resolve the issue.



## Charge Back & Accounting

Cloudian HyperStore provides system-wide, detailed reporting of usage and cost based for a group or individual user rate of consumption such as storage size, data transfer (in and out), number of object get and put operations as well as number of metadata get and put operations. The first step is to define a rating plan. This is the phase where a price is associated to a particular metric; for example, for consumption of storage size up to 100GB, it would be \$.1/GB per mo. Automatically, statistics are recorded about the users and groups' consumption; finally, at the end of the month, end of billing cycle, a report can be generated calculating the bill for each of groups or users based on usage rates.



## Tunable Consistency Levels

HyperStore also provides the ability to configure the level of data consistency when using replication or EC to protect objects across sites. For example, the default consistency requirement for read and write operations can be defined as "quorum", meaning that a read or write operation must succeed on a quorum (or majority) before a success response is returned to the client application. For those data objects that are considered mission critical, the system can wait until acknowledgment is received from ALL nodes across a single or multiple datacenter locations.



## One Platform for Objects & Files

Cloudian HyperStore Connect for Files (CHCF) allows file services on top of Cloudian HyperStore object storage using industry standard protocols like NFS, CIFS & FTP. File-based environments can gain tremendous benefits from the rich feature set intrinsic to Cloudian HyperStore object storage system such as high durability and availability, geo dispersal, multi-tenancy and low cost.



## Object Lifecycle Management

With Cloudian HyperStore, you can configure expiration for objects through Object Lifecycle Management. For example, you can impose a policy on an object so that after a year, the object will be deleted (expired) automatically. Through object publishing, you can make an object available and accessible via the web. You can publish it through a URL, allow a number of downloads against the object and place an expiration. URL access will vanish accordingly.



## Storage Optimization

With Cloudian HyperStore, first and foremost, objects determined to be small in size (for example, objects less than 128KB) can be stored directly in Cassandra as a blob. The advantage is that these Objects stored in Cassandra will not incur the HyperStore File system overhead. Secondly, Cloudian HyperStore provides three types of compression modules available to choose from, namely - LZ4, snappy and Zlib. Pre-hand knowledge of the dataset helps determine the best possible compression module to implement for the objects.