

The Lenovo logo is displayed in white text on a black rectangular background.

Reference Architecture: Lenovo Storage DX8200C Powered by Clouidian

Last update: 6 June 2016

Shipped from the factory as a fully integrated appliance

Can be used to build a public or private cloud

Provides a highly scalable multi-tenant storage system

100% compliant with Amazon S3 API

Mike Perks, Lenovo

Bill Hansen, Lenovo

Phillip Tribble, Clouidian



Table of Contents

1	Introduction	1
1.1	Executive summary	1
1.2	Audience	2
2	Business problem and business value	3
2.1	Business problem	3
2.2	Business value	4
3	Requirements	5
3.1	Functional requirements	5
3.2	Non-functional requirements	5
4	Architectural overview	6
4.1	Cloudian HyperStore architecture	6
4.2	S3 API compatibility	6
5	Component model	8
6	Operational model	10
6.1	Hardware components	10
6.1.1	Storage server	10
6.1.2	Network switches	11
6.1.3	Lenovo XClarity	12
6.2	Storage servers	14
6.2.1	Storage server capacity planning	14
6.3	Networking	16
6.4	Data architecture	17
7	Deployment considerations	18
7.1	High availability	18
7.2	Cloudian Management Console	18
7.3	Load balancers	18

7.4	Backup.....	18
7.5	Tiering.....	18
7.6	Multi-tenancy.....	18
8	Appendix: Bill of Materials.....	19
8.1	Part Number for Appliance.....	19
8.2	BOM for networking.....	20
	Resources	21

1 Introduction

This document describes a reference architecture for the Lenovo Storage DX8200C powered by Cloudian and is a result of a collaboration between Lenovo and Cloudian.

1.1 Executive summary

With the massive growth of data from the Internet of Things (IOT) to collaboration to compliance, users are demanding low-cost, flexible, easy to scale, and simple to manage data center storage solutions.

Software-defined object storage delivers on these demands by capitalizing on industry standard x86 infrastructure and storage technologies to deploy more economic and manageable storage solutions compared to legacy storage architectures.

The Lenovo Storage DX8200C powered by Cloudian is an example of the new breed of software-designed storage. This appliance allows companies to build their own public or private cloud storage infrastructure including enterprise IT organizations, cloud service providers, or cloud hosting providers. The decision to use cloud storage for the delivery of IT services is best made by starting with the knowledge and experience that is gained from previous work. This document gathers the essential information about a scale-out storage reference architecture and a real-world example from the Cloudian support organization that uses the Lenovo appliances that are powered by Cloudian software running on Lenovo's award winning hardware.

Enterprise and managed solution providers who are looking to deploy enterprise file sync and share software can use the Lenovo appliance's object storage to store and access user files from any location and any device of their choice. The DX8200C also provides a central trusted backup and file share repository and sharing server for regional offices. Enterprises can now run Hadoop analytics directly on the Lenovo Storage DX8200C powered by Cloudian. This in-place analytics enables customers to derive meaningful business intelligence from their data quickly, efficiently, and economically.

The Lenovo Storage DX8200C provides administrators with the ability to deploy data in different data centers or geographically separated regions. The Lenovo appliance also connects different clouds with the ability to tier data to Amazon S3, Amazon Glacier, or another HyperStore deployment. With multi-tenant support, an organization can provide storage for multiple departments or groups of users without the need to deploy additional infrastructure. It also includes necessary features such as multi-currency rating plans, billing, quality of service metrics, enforcement, and an API that allows service providers to integrate Cloudian into their own portal.

The combination of Cloudian HyperStore and Lenovo world-class hardware provides a unique 100% S3 compliant multi-device cloud storage platform with which any level of SLA can be met at the right cost with dynamic migrations from one storage tier to another. The solution meets the needs of any enterprise willing to adapt their storage cost to the value of their data as well as service providers targeting to market multi-SLA storage services for their customers.

The Lenovo Storage DX8200C powered by Cloudian solution is part of the Lenovo StorSelect Software Defined Storage (SDS) program, which provides quality, engineering prowess, and Lenovo's worldwide global services. Customers can quickly and confidently deploy new software-defined storage technologies, using pre-loaded appliances with software from best-of-breed Independent Software Vendors (ISVs) on proven Lenovo systems. The StorSelect solutions enable easy expansion for scalable deployments, as well as

integration with existing storage infrastructure, so that you can grow your SDS deployment on your own terms.

1.2 Audience


The intended audience of this document is IT professionals, technical architects, sales engineers, and consultants to assist in planning and deploying the Lenovo Storage DX8200C powered by Cloudian appliance. This paper also assumes that the reader has basic knowledge of cloud storage infrastructure components and services.

2 Business problem and business value

This section provides a summary of the business problems that this reference architecture is intended to help address, and the additional value that this solution provides for storing large amounts of data.

2.1 Business problem

An organization's data continues to grow each year. According to analysts (such as the IDC), data is doubling in size every two years (see Figure 1). The yearly storage increase creates a need for large capacity storage. It is crucial to have a storage solution ready to scale to the needs of its users.



- **Estimated 1.8 Zettabytes of data in the world in 2010 by IDC**
 - Fun fact “1 Zettabyte = 31.6 Billion iPads 32GB”

Source: IDC, 2014

Figure 1. Predicted data growth

This solution features wide ranging applicability and is particularly strong for the following use cases:

- Backup and archival for big data

As electronic data has exploded, backup windows have shrunk. Despite this new era of big data, IT and backup administrators are asked to use antiquated methods to manage new world demands. As a result, IT and backup administrators are spending an inordinate amount of time and budget managing onsite and offsite backups and reactively dealing with data recoveries as a fire drill. Working with backup software partners,.

- Smart data analytics

Today's smart businesses are looking at ever increasing use of data analytics to derive business value and provide a better customer experience.

- Remote Office Backup

Many Enterprises today are geographically dispersed and have many remote offices. These remote offices must share and collaborate on content and the data must be protected and backed up. Enterprises want to enhance sharing across these remote locations, protect and secure the business data, and do this with central IT administration and control.

- Enterprise file synchronization and sharing

With the proliferation of mobile devices and the digitization of content, there is a need for new solutions to enhance collaboration between users in Enterprises. Traditional NAS/SAN infrastructures do not support mobile devices well and do not support the distributed nature of the content.

- Private and hybrid cloud storage

Modern Enterprises are evolving their IT organizations into service-driven IT that implement services on demand, such as the Public Cloud. Organizations are implementing S3 Private Clouds, Cloud Platforms (such as CloudStack and OpenStack) to implement this Private Cloud environments.

- Service Provider-as-a-Service

Traditionally, service providers used storage technologies (specifically, Network Attached Storage [NAS] and Storage Area Networks [SAN]) to offer storage-as-a-service. However, these traditional storage technologies are better suited for small enterprise deployments and do not properly scale to handle large cloud deployments.

2.2 Business value

To solve this massive growth of data storage problem, organizations are looking for new approaches to storage. These modern solutions typically have the following key features:

- Scale out architecture
- Low-cost commodity hardware
- Hybrid Cloud Tiering
- S3 compatible
- Multiple data centers
- Multi-tenancy

In addition, many data center administrators need the ease of deploying a “turn-key” storage offering to scale out their storage infrastructure. For these organizations, there is value to implementing pre-certified storage appliances that take the guess work out of configuring the right server and storage combinations. Rapid deployment and risk mitigation are key benefits for these enterprises, especially because their IT staffs are stretched too thin. A turn-key system helps to speed up deployment times and mitigate the risk of application downtime or performance problems that can occur from misconfigured systems. Another benefit is that it gives the IT organization a single vendor to support the entire hardware and software stack

3 Requirements

This section describes the functional and non-functional requirements for this solution.

3.1 Functional requirements

Table 1 lists the functional requirements for an Enterprise cloud object store.

Table 1: Functional Requirements

Requirement	Description
Multi-tenancy	Provide storage for different customers, groups, or organizations
Quality of Service (QoS)	Set QoS settings for groups, tenants, user, or entire system
Multi-region and data center	Store data in geographically dispersed regions or data centers
Billing and reporting	Generate bills for users and integrate into external billing systems
Encryption	Server-side encryption of files
Monitoring	JMX and SNMP charge back reports
Access Control Lists (ACL)	Set ACL's for a bucket or object
Bucket protection	Choose different protection levels, such as replication or erasure code

3.2 Non-functional requirements

Table 2 lists the non-functional requirements for an Enterprise cloud object store.

Table 2: Non-functional Requirements

Requirement	Description
Scalability	Solution components scale for growth
Load balancing	Workload is distributed evenly across servers
Fault tolerance	Single component error does not lead to whole system unavailability
Physical footprint	Compact solution
Ease of installation	Reduced complexity for solution deployment
Ease of management/operations	Reduced complexity for solution management
Flexibility	Solution supports variable deployment methodologies
Security	Solution provides means to secure customer infrastructure
Power efficient	Uses low-power CPU's and power efficient drives with large capacities

4 Architectural overview

This section presents an architectural overview of the Lenovo Storage DX8200C powered by Clouidian solution and compatibility with the Amazon S3 API.

4.1 Clouidian HyperStore architecture

Clouidian HyperStore® software uses a fully distributed and replicated peer-to-peer architecture with no single point of failure. It easily scales horizontally by using commodity hardware so deployments can start with a few servers in a single data center and then scale out as usage increases to thousands of servers that are distributed across multiple data centers that are managing hundreds of petabytes of data. Its distributed architecture with automatic replication and recovery services makes it highly resilient to network and node failures without data loss. Similarly, when scaling the storage cluster or performing maintenance, changes in node availability are automatically detected without service interruption. Figure 2 shows an overview of Clouidian HyperStore with the DX8200C.

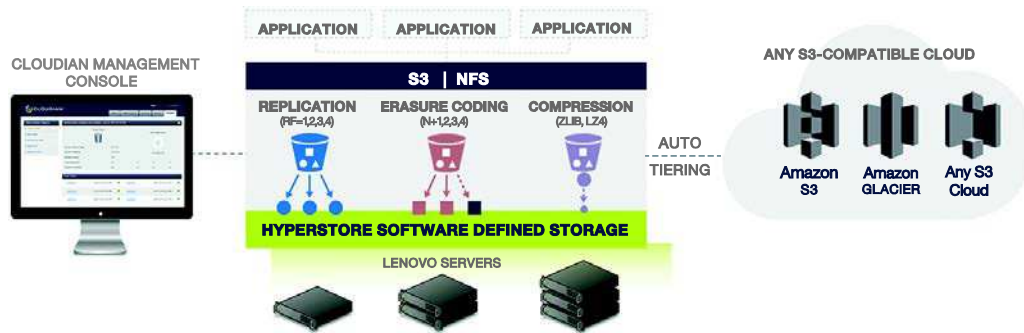


Figure 2. Lenovo Storage DX8200C powered by Clouidian architecture

Clouidian HyperStore data can be stored in erasure code or a replicated format. Replication results in greater performance, but less storage capacity. Erasure code provides the most cost-effective way to store data because the data is stored in smaller chunks across the cluster. However, erasure code requires more system resources, which leads to greater latency on storage requests. Replication uses the most storage space because by default the replication factor is 3, which means that the object is copied to three different nodes.

Clouidian HyperStore supports multiple compression algorithms that compresses objects on the storage side to save space. Choosing the right storage type depends on your storage size and needs. If the media content stored on the cluster is meant for archiving and not accessed frequently, erasure code is the best choice. If the media content requires the greatest performance and replication size is not a factor, replication is the best choice.

4.2 S3 API compatibility

Amazon S3 commands twice the market share of all its closest competitors combined and it likely is the storage platform of choice for on-premise hybrid or private cloud deployments. Companies and developers that are implementing applications with the Amazon S3 API depend significantly on its compatibility from service providers to ensure that their applications remain compatible and function. With no standards enforced for

claiming S3 compatibility, choosing the right storage platform can be tenuous.

The Lenovo Storage DX8200C storage platform offers 100% advanced compatibility and allows developers continued use of Amazon's S3 SDK (Software Development Kit). By supporting native S3 API calls, developers can significantly ease their workloads by not changing SDK's or API's. Also, the DX8200C can automatically tier data between on-premise cloud deployments and Amazon's S3 public cloud while representing the cloud ecosystem under a single name space. With this advanced functionality, the Lenovo appliance provides a compatible storage platform for S3 on-premise and hybrid cloud deployments.

Figure 3 shows a comparison of Clouidian HyperStore and Amazon S3 object storage.

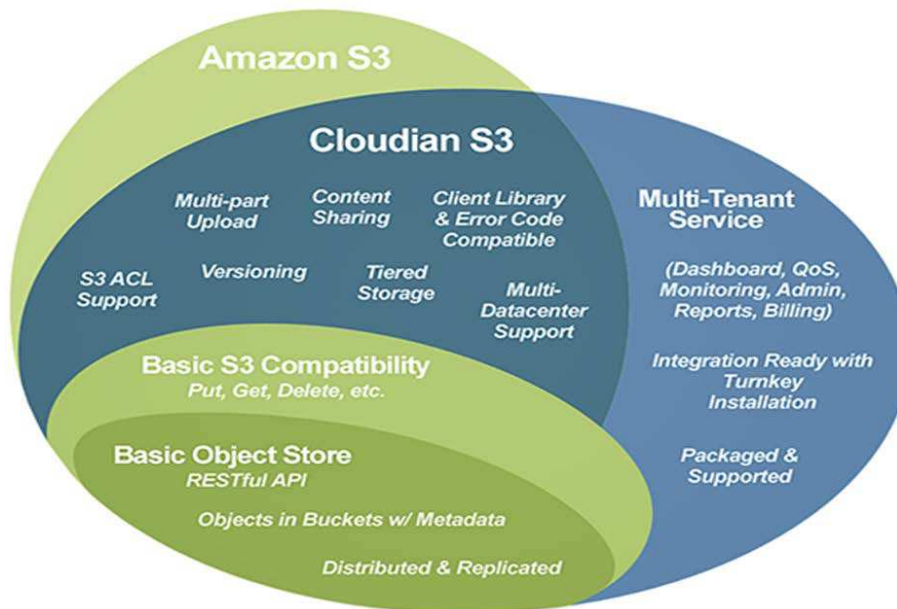


Figure 3. Clouidian HyperStore S3 environment

5 Component model

The Lenovo Storage DX8200C powered by Clouidn offers flexibility with the software installer to provide the best source of flexibility for an environment. There are standard components, such as the S3 server, HyperStore server, Clouidn Management Console, Cassandra, Redis, and Puppet that make up and define the Clouidn HyperStore architecture, as shown in Figure 4. This section describes each component that is needed for the installation and data architecture.

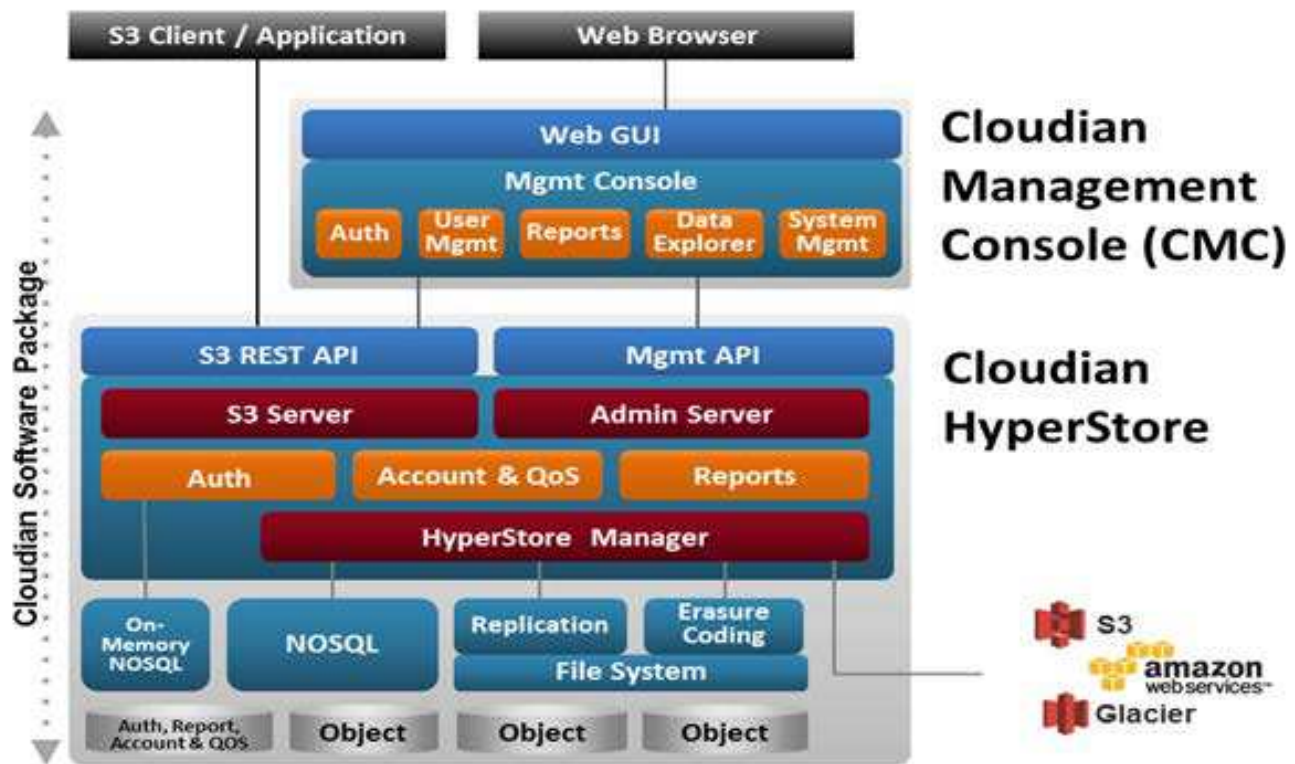


Figure 4. The Clouidn HyperStore component model

The Clouidn HyperStore software features the following main components:

Clouidn Management Console (CMC)	The CMC is a web-based interface for Clouidn HyperStore system administrators, group administrators, and users. From the console, administrators can manage users, generate usage reports and bills, view data that is stored on the system, and access system management functions. The system management functions allow administrators to setup log alerts and email notifications, view the storage utilization on the cluster or individually on the nodes, and view the status of the HyperStore services. In addition, node management offers the ability to add or remove a node, perform data consistency checks on the nodes, and view meta data information for objects.
----------------------------------	---

S3 Service	<p>The HyperStore software provides a high-performance S3 proxy service. The S3 Service processes S3 REST requests incoming from client. The S3 service handles the initial login request from the S3 client with the access and secret key. After logging in, Account and QoS settings are checked for each request. For example, if there is a QoS rule for a user to have a maximum of 100 MB storage and the user tries to upload 200 MB, the request is unsuccessful. The S3 daemon also contains the Admin service that extends the S3 API by allowing administrators to manage users and groups, S3 credentials, billing, and generate usage reports.</p>
HyperStore service	<p>As an object store, Cassandra provides a wealth of valuable built-in functionality including data partitioning, automatic replication, easy cluster expansion, quorum calculation, and so on. After the S3 service honors the request, HyperStore must access and manage the data. For example, a PUT object request comes to the S3 service and then a request is sent to HyperStore to store the data. Depending on the storage protection type, files might be stored in erasure code, Cassandra file system, or replication.</p>
NoSQL	<p>The HyperStore software uses the open source storage platform Cassandra to store several types of data, including object meta data, user information, and accounting information. S3 client applications do not access Cassandra databases directly; all S3 access is mediated through the S3 Service. The HyperStore Service and Admin Service also access Cassandra.</p>
Redis	<p>The HyperStore software uses the lightweight, open source Redis key-value data store to store various data that supports HyperStore S3 service features. The Redis Credentials DB stores user credentials and other S3 operation supporting data, such as multi-part upload session information and public URL access counters. The Redis QoS DB stores user-level and group-level Quality of Service settings that were established by system administrators. The DB is also used to keep count of user requests so that Quality of Service limits can be enforced by the system.</p>

6 Operational model

This section describes the validated Lenovo Storage DX8200 powered by Cloudian .. It concludes with some example deployment models.

6.1 Hardware components

The Cloudian HyperStore software was validated to run on the Lenovo servers and Lenovo RackSwitch network switches.

6.1.1 Storage server

The DX8200C appliance is comprised of the proven Lenovo System x3650 M5 servers running the Cloudian HyperStore software.

Lenovo Storage DX8200C powered by Cloudian

The Lenovo Storage DX8200C powered by Cloudian appliance is part of the award-winning Lenovo portfolio of systems. This versatile powerhouse incorporates industry-leading reliability, security, energy efficiency and world-class performance for business-critical applications and cloud deployments. Lenovo servers ranked #1 in reliability of all x86 servers for the third year in a row and #1 in overall customer satisfaction for 2014 and 2015. The DX8200C offers a flexible, scalable design and simple upgrade path to fourteen 3.5-inch hard disk drives (HDDs) for up to an impressive 112TB of storage capacity per node.

Lenovo XClarity provides centralized management including automated discovery, inventory tracking, real-time monitoring, configuration, fault detection and alert handling.

For easy network expansion and flexibility the server's onboard Ethernet solution provides two standard embedded Gigabit Ethernet ports and two 10 Gigabit Ethernet ports without occupying PCIe slots. The appliance also incorporates smart energy features for data center savings, including 80 PLUS® Titanium power supplies for up to 96 percent efficiency and TruDDR4 Memory using 45 percent less energy than DDR3 memory.

Lenovo systems accelerate performance across virtualized, transactional databases, high performance computing, and high frequency trading applications. Combined with the powerful Intel® Xeon® processor E5-2600 v4 product family, the Lenovo system offers outstanding workload density, designed to lower the total cost of ownership (TCO) per virtual machine.



Figure 5. Lenovo Storage DX8200C powered by Cloudian

6.1.2 Network switches

The Lenovo Storage DX8200C powered by Cloudian can be used with the Lenovo RackSwitch network switches that are described in this section.

Lenovo RackSwitch G8124E

The Lenovo RackSwitch G8124E (as shown in Figure 6) delivers exceptional performance that is lossless and low-latency. It also provides high availability and reliability with redundant power supplies and fans as standard. In addition, RackSwitch G8124E delivers excellent cost savings and a feature-rich design regarding virtualization, Converged Enhanced Ethernet (CEE)/Fibre Channel over Ethernet (FCoE), Internet Small Computer System Interface (iSCSI), and enterprise-class Layer 2 and Layer 3 functionality.

With support for 10 Gb, this 24-port switch is designed for clients who are using 10 Gb Ethernet or plan to do so. The G8124E is designed to support Lenovo Virtual Fabric, which provides the ability to dynamically allocate bandwidth per virtual network interface card (vNIC) in increments of 100 MB, while adjusting over time without downtime.



Figure 6. Lenovo RackSwitch G8124E

For more information, see the [RackSwitch G8124E Product Guide](#).

Lenovo RackSwitch G8264

Designed with top performance in mind, Lenovo RackSwitch G8264 (as shown in Figure 7) is ideal for today's big data, cloud, and optimized workloads. The G8264 switch offers up to 64 10 Gb SFP+ ports in a 1U form factor and allows for future expansion with four 40 Gb QSFP+ ports. It is an enterprise-class and full-featured data center switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center grade buffers keep traffic moving. Redundant power and fans and numerous high-availability features equip the switches for business-sensitive traffic.




Figure 7. Lenovo RackSwitch G8264

The G8264 switch is ideal for latency-sensitive applications, such as client virtualization. It supports Virtual Fabric to help clients reduce the number of I/O adapters to a single dual-port 10 Gb adapter, which helps reduce cost and complexity. The G8264 switch supports the newest protocols, including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for support of FCoE, iSCSI, and NAS.

For more information, see the [RackSwitch G8264 Product Guide](#).

Lenovo RackSwitch G7028

The Lenovo RackSwitch G7028 (as shown in Figure 8) is a 1 Gb top-of-rack switch that delivers line-rate Layer

2 performance at an attractive price. G7028 has 24 10/100/1000BASE-T RJ45 ports and four 10 Gb Ethernet SFP+ ports. It typically uses only 45 W of power, which improves energy efficiency.



Figure 8. Lenovo RackSwitch G7028

For more information, see the [RackSwitch G7028 Product Guide](#).

Lenovo RackSwitch G8052

The Lenovo System Networking RackSwitch G8052 (as shown in Figure 9) is an Ethernet switch that is designed for the data center and provides a virtualized, cooler, and simpler network solution. The Lenovo RackSwitch G8052 offers up to 48 1 GbE ports and up to four 10 GbE ports in a 1U footprint.



Figure 9. Lenovo RackSwitch G8052

For more information, see the [RackSwitch G8052 Product Guide](#).

6.1.3 Lenovo XClarity

Lenovo XClarity™ Administrator is a centralized resource management solution that reduces complexity, speeds up response, and enhances the availability of Lenovo® server systems and solutions.

The Lenovo XClarity Administrator provides agent-free hardware management for Lenovo's System x® rack servers and Flex System™ compute nodes and components, including the Chassis Management Module (CMM) and Flex System I/O modules. Figure 10 shows the Lenovo XClarity administrator interface, where Flex System components and rack servers are managed and are seen on the dashboard. Lenovo XClarity Administrator is a virtual appliance that is quickly imported into a virtualized environment server configuration.

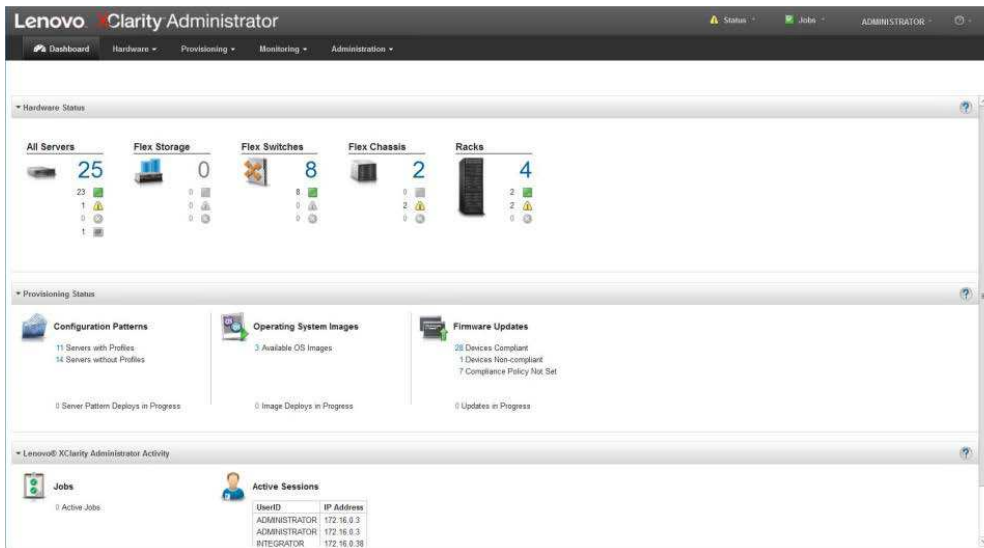


Figure 10: XClarity Administrator interface

Lenovo XClarity Administrator is an optional software component for the Lenovo Storage DX8200C powered by Cloudian product, which can be used to manage firmware upgrades outside of the Cloudian web console.

6.2 Storage servers

This section describes the hardware configuration and capacity planning examples for a typical Lenovo Storage DX8200C powered by Cloudian deployment. Table 3 lists the default configuration for the storage servers. The table below represents the parts for a single node. The DX8200C configuration requires a minimum of three nodes – with a total base capacity of 252TB (using 14 x 6TB drives per node) or 336TB (using 14 x 8TB drives per node). Customers can order additional nodes as needed to increase capacity.

Table 3: Default Hardware Configuration for Lenovo Storage DX8200C

	Lenovo Storage DX8200C (Single Node)
Form Factor	2U
Drives	14 x 6 TB or 14 x 8TB
Storage Capacity	84 TB or 112TB
CPU	Single Intel Xeon E5-2630 v4 Series processor 8C 2.4GHz
Memory	64 GB
3.5" Disks Installed	14 x 6TB 7.2k SATA or 14 x 8TB 7.2k SATA
2.5" Disks installed	2 x 240GB SATA SSD (for 84TB base) or 2x480GB SATA SSD (for 112TB base)
Network Interface	10 GbE dual port NIC10 GbE dual port NIC
RAID Adapter (pass-through)	N2215 and M1215
Power Supplies	2 x 750W
Data Protection	Replication and erasure coding
Protocol Support	S3, NFS
Monitoring/Management	CLI, GUI, API, IPMI, JMX
Cloudian HyperStore License + 3 year support	Included

The following options can be added to the base model:

- Memory
- Rails/Slides
- Ethernet adapters
- Quantity of HBA controllers
- Quantity of power supplies

6.2.1 Storage server capacity planning

The Lenovo Storage DX8200C supports data replication and erasure coding, which have their optimal use cases. It is important to choose an optimal protection approach for your data. The most important trade-off is the capacity efficiency that is given by data replication versus the reduced I/O latency that is given by erasure coding.

For data replication, the minimum number of servers is three. To estimate how much usable data is available in the cluster when replication is used, take the total raw capacity of the disks in the HyperStore cluster and divide by three.

Table 4 lists the number of DX8200C units that are required, assuming 14 x 6TB drives in each server.

Table 4: HyperStore Replication sizing (minimum 3node start – 3way Replication)

Usable Data	84TB	168 TB	252 TB	336 TB
Units Required	3	6	9	12

For erasure coding, the default number of units is six: four units for data fragments and two units for redundant data fragments. To estimate the total usable space for erasure code, take the total raw capacity of the disks in the Lenovo Storage DX8200C powered by Cloudian cluster and divide by the sum of the data and redundant fragments that is specified in the erasure code configuration.

Table 5 lists the number of servers required assuming 14 x 6TB drives in each server.

Table 5: HyperStore Erasure coding sizing (minimum 6node start)

Usable Data	Min 6node	336 TB	504 TB	672 TB
Units Required	3	6	9	12

6.3 Networking

This section describes the networking topology and includes design guidance to correctly configure the network environment for redundancy and failover.

This reference architecture uses two 24-port ultra low-latency, high-performance Lenovo RackSwitch G8124 10 GbE network switches to provide primary data communication services. If more ports are required, the 64-port Lenovo RackSwitch G8264 switch can be used. The management interface on the compute nodes can also be connected to one of the two switches or the management interface can connect to an extra 1 GbE switch, such as the Lenovo RackSwitch G7028.

High availability and failover in the network architecture is achieved by using Inter-Switch Link (ISL), Link Aggregation Control Protocol (LACP), and Virtual Link Aggregation Groups (vLAGs). The recommended vLAG/LACP configuration is shown in Figure 11.




Figure 11. LACP/vLAG recommended network design

An ISL is a physical network connection from a physical network port on one switch to a physical network port on another switch that enables communication between the two switches. This reference architecture uses two physical connections between the two networking switches, which are link aggregated.

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (which is known as a *dynamic trunk group*). If a link in a LACP trunk group fails, traffic is reassigned dynamically to the remaining links of the dynamic trunk group.

LACP teams are formed on the ISLs between the switches and on the host connections to the switches, which provides for host connection redundancy. To maintain maximum bandwidth over the multiple connections, vLAGs also are configured on the LACP teams. Disabling Spanning Tree on the LACP teams helps avoid the wasted bandwidth that is associated with links that are blocked by spanning trees.

By using VLAGs, the redundant uplinks remain active and use all available bandwidth. To maintain maximum bandwidth over the multiple connections, vLAG is enabled on the LACP teams in this reference architecture.

At the operating system level, the NIC ports are bonded together to provide high availability and failover for the Cloudian HyperStore.

6.4 Data architecture

Each Lenovo Storage DX8200C powered by Cloudian implementation starts with three or more distributed nodes and then objects are replicated or erasure coded across the available nodes for data durability and availability. Administrators can configure the number of replicas or erasure code strategy required to meet SLA and cost objectives, including the option to replicate copies to other data centers for geo redundancy. Reads and writes are always performed at the local data center with remote replication performed in the background to avoid latency of remote writes.

Figure 12 shows an example of a Lenovo Storage DX8200C powered by Cloudian cluster that is distributed and elastic across Geos.

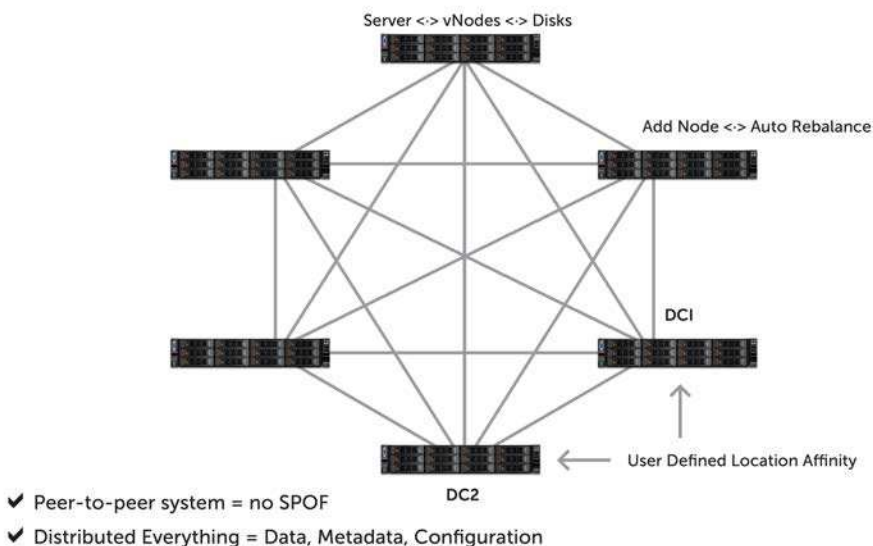


Figure 12. Lenovo Storage DX8200C powered by Cloudian data architecture

7 Deployment considerations

This section describes several considerations, such as high availability, Server/CMC nodes, load balancers, and backup, and tiering to different cloud infrastructures.

7.1 High availability

The Lenovo Storage DX8200C powered by Cloudbian appliance contains high availability attributes. By default, every node includes the core services installed and running, such as Cassandra, Redis, HyperStore, S3, and the Cloudbian Management Console. Redis services are installed in a master and slave and can fail over to another node if a failure is detected. For example, a deployment of six nodes can tolerate up to two node failures if the default read and write consistency level of quorum is met.

7.2 Cloudbian Management Console

The Cloudbian Management Console (CMC) can be installed independently on a separate node in the cluster that is running core services, such as Cassandra, S3, HyperStore, and Redis. The CMC also can work in a load balancer configuration for high-availability and to scale performance. Alternatively, the CMC can be run on a less powerful server separate from the data nodes.

7.3 Load balancers

The Cloudbian HyperStore software works with most load balancers that are available on the market. Services that are available for load balancing are: S3, Admin, and the CMC.

7.4 Backup

By using third-party applications, data that is stored in Cloudbian HyperStore can be backed up to other storage mechanisms, such as tape, another Cloudbian HyperStore deployment, or other cloud storage providers.

7.5 Tiering

Cloudbian HyperStore allows users to transition objects on a per-bucket level basis to Amazon S3, Amazon Glacier, or another Cloudbian HyperStore cluster.

7.6 Multi-tenancy

Multi-tenant support enables service providers to offer storage to multiple customers. System administrators can isolate different customers or departments of an organization into separate groups on a single HyperStore cluster. Each tenant can have group administrators to manage administration tasks such as adding or removing authorized users, maintaining and implementing quality of service settings, and creating usage reports.

8 Appendix: Bill of Materials

This appendix contains the Bill of Materials (BOMs) for different configurations of hardware for Cloudian HyperStore deployments. There are sections for storage servers and networking.

8.1 Part Number for Appliance

Below is the part number for a single Lenovo DX8200C appliance.

Part #	Description	Quantity
5120A1C	Lenovo Storage DX8200C	1

8.2 BOM for networking

Below is the bill of materials for a single 1 GbE network switch.

Part #	Description	Quantity
7159BAX	Lenovo RackSwitch G7028 (Rear to Front)	1
39Y7938	2.8m, 10A/100-250V, C13 to IEC 320-C20 Rack Power Cable	2

Below is the bill of materials for a single 10 GbE network switch.

Part #	Description	Quantity
7159BR6	Lenovo RackSwitch G8124E (Rear to Front)	2
39Y7938	2.8m, 10A/100-250V, C13 to IEC 320-C20 Rack Power Cable	4
90Y9427	1m Passive DAC SFP+ Cable	2

Resources

For more information, see the following resources:

- DX8200C powered by Clodian web page:
<http://shop.lenovo.com/us/en/systems/storage/sds/cloudian>
- DX8200C powered by Clodian data sheet:
http://www.lenovo.com/images/products/system-x/pdfs/datasheets/lenovo_storselect_dx8200c_ds.pdf
- StorSelect 3D tour:
<http://www.lenovofiles.com/3dtours/products/superblaze/storselect/index.html>
- StorSelect video:
<http://www.lenovofiles.com/videos/storselect.html>
- StorSelect analyst paper:
http://www.lenovo.com/images/products/system-x/pdfs/analyst-reports/storselect_ar.pdf
- Clodian HyperStore resources:
<http://cloudian.com/resources>
- Clodian HyperStore Technical Review white paper:
<http://cloudian.com/cloudian-docs/white-papers/Cloudian-HyperStore-Technical-Review.pdf>
- S3 API: What it takes to build 100% Compatible Object:
<http://info.cloudian.com/S3API-wp-lp-0515.html>
- Using Clodian HyperStore with Hortonworks:
<http://cloudian.com/cloudian-docs/Cloudian-Hortonworks-Solutions-Brief.pdf>

Trademarks and special notices

© Copyright Lenovo 2016.

© Copyright Clouidian Inc 2016

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

Clouidian, the Clouidian logo, and HyperStore® are registered trademarks or trademarks of Clouidian, Inc. All other trademarks are property of their respective holders.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.