



Enterprise Data Protection with Hyperconverged Infrastructure and Object Storage

Describes hyperconverged solutions for data backup, archive and disaster recovery

Introduces a solution using Lenovo, Cloudian, Rubrik and Nutanix components

Describes Archival and Disaster Recovery use cases

Architectural and implementation guidance for IT Specialists, Solution Architects and Administrators

Grant Jacobsen
Bill Hansen



Abstract

This paper describes the architecture, implementation and benefits of an integrated solution that combines three key technologies:

- ▶ Lenovo® Storage DX8200C powered by Cloudian
- ▶ Lenovo ThinkAgile™ HX series powered by Nutanix
- ▶ Rubrik Cloud Data Management

This solution provides an end-to-end software defined infrastructure that does the following:

- ▶ Eases maintenance and lowers total cost of ownership
- ▶ Scales into the petabytes
- ▶ Provides hyperconverged simplicity and performance
- ▶ Delivers best-in-class integrated data management and backup capabilities
- ▶ Handles rapid business growth and has the flexibility to support evolving compute and storage needs

This paper is intended to provide IT professionals, technical specialists, solutions architects and administrators with information about the integrated solution. The paper assumes the reader is familiar with software defined storage, hyperconvergence, and cloud storage infrastructure components and services.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges. See a list of our most recent publications at our web site:

<http://lenovopress.com>

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction	3
Components	3
Solution Overview	8
Solution testing and validation	13
Use Cases	15
Conclusion	19
Authors	19
Notices	20
Trademarks	21

Introduction

Data center modernization is now a business imperative. Enterprises of all sizes must balance current customer demands, rapid business growth, and new customer requirements while implementing innovations that will drive new revenue streams. The ability to quickly adapt to change can be a challenge with traditional infrastructures.

At the same time, data growth is exploding with the amount doubling roughly every two years. According to the Economist Intelligence Unit, it is estimated that 90% of the world's data was created in the past two years, with 80% of that data unstructured. In addition, internal and external data regulations are increasingly requiring instant data accessibility, retention governance and compliance requirements.

This dramatic growth in data increases the need for low cost, flexible, and easy to scale data management, protection and storage solutions. Soon, it will not be uncommon for enterprises to have Petabytes of data. The ability to manage data at this scale using traditional infrastructure is not easy and can be very expensive. High performance and availability are also difficult at this scale. With storage lifecycle expectations increasing, the challenge to keep storage affordable at petabyte scale is even more acute.

Today, there is a solution that solves these issues and delivers on these requirements:

- ▶ Lenovo DX8200C powered by Cloudian appliance
- ▶ Lenovo ThinkAgile HX Series appliance powered by Nutanix
- ▶ Rubrik Cloud Data Management (CDM) Appliance

These products combine to provide an end-to-end software defined infrastructure all enabled by Lenovo. With this solution, Enterprises and Service Providers benefit from a highly scalable hyperconverged cluster, petabyte-scale object storage, and data management and protection capability.

Components

This section introduces the major components of the solution:

- ▶ “Lenovo Storage DX8200C powered by Cloudian”
- ▶ “Cloudian HyperStore” on page 4
- ▶ “Lenovo ThinkAgile HX Series powered by Nutanix” on page 5
- ▶ “Nutanix” on page 6
- ▶ “Rubrik Cloud Data Management” on page 7

Lenovo Storage DX8200C powered by Cloudian

The Lenovo Storage DX8200C is built on the industry-leading Lenovo x3650 M5 server that features enterprise-class reliability, management, security, and world-class performance.

Pre-integrated and pre-validated with Cloudian HyperStore software, this appliance provides affordable, highly efficient, high capacity storage to address the rapid growth of unstructured data. Unstructured data includes media images, video, web content, and file shares.



Figure 1 Lenovo Storage DX8200C powered by Cloudian

The DX8200C features a flexible, scalable design with fourteen 3.5-inch hard disk drives (HDDs) for up to an impressive 168TB of storage capacity per node. The DX8200C is available in 5 different models, with capacities of 56TB, 84TB, 112TB, 140TB, and 168TB. A base configuration requires a minimum of three nodes. As data storage requirements increase, users can expand and scale the solution by simply adding additional DX8200C nodes – up to hundreds of Petabytes and with a single, global namespace¹.

For more information, see the Lenovo Storage DX8200C reference architecture:

<https://lenovopress.com/lp0531-lenovo-storage-dx8200c-powered-by-cloudian-ra>

Cloudian HyperStore

Cloudian HyperStore software is pre-installed on the Lenovo DX8200C. HyperStore is designed for enterprises and service providers with data-intensive environments where unstructured data grows exponentially on an annual basis straining the limits of the traditional data center. Cloudian HyperStore provides infinite scalability, data protection with up to 14 nines of data durability, and powerful data management - at a monthly cost as low as \$0.01 per GB². HyperStore customers benefit from a rich set of service provider-oriented features including quality of service controls, multi-tenancy, and billing. HyperStore also includes 100% native S3 API compatibility which ensures plug-and-play interoperability with S3-enabled applications.

Key capabilities and features of Cloudian HyperStore:

- ▶ **Infinite scalability** on demand allows enterprises to seamlessly extend their storage capacities based their need, from as little at 56TB, to hundreds of Petabytes, across multiple data centers to meet increasing demand.
- ▶ **100% Native S3 compatibility** offers the flexibility to be an on-premises object storage as well as a hybrid tier to Amazon, Google and other public clouds.
- ▶ **Hybrid Cloud Backup and Archive** combines a petabyte-scalable, high-performance on-premises backup target, seamless cloud storage integration, and full interoperability with Rubrik data management and protection. Together, this hybrid cloud implementation lets you retain a familiar workflow while ensuring successful backup window predictability and repeatable Recovery Time Objective/Recovery Point Objective (RTO/RPO).
- ▶ **Multi-tenancy** of the HyperStore system provides the capability to securely have multiple users reside on a single, shared infrastructure. 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.

¹ Namespace: enables a single, unified view of the file system data across all physical storage locations

² \$0.01 per GB per month – based on total hardware plus software cost amortized over 3 years

- ▶ **Auto-tiering** within HyperStore allows for data to automatically move from a Cloudian cluster to S3 compliant storage, as well as HyperStore to HyperStore, based on a configurable schedule.
- ▶ **Quality of Service (QoS)** allows the definition of specific classes of service for groups and users and configuration billing and charge-back policies. Both administrators and users benefit from unique reporting operations and account and data management capabilities.
- ▶ **Integrated billing, management and monitoring system** maintains comprehensive service usage data for each group and each user in the system.

Central to Cloudian's data protection are storage policies allowing users to set protection to meet their targeted level of durability and availability. The policies are built of a flexible design using replication, erasure coding and encryption:

- ▶ **Erasure Coding** enables deep archive efficiency and flexible redundancy, providing robust data protection without consuming precious disk space. Object replicas are employed for frequently used data.
- ▶ **Replication** allows for a configurable number of copies of each data object to be maintained in the system, with each copy stored on a different node.
- ▶ **Server-Side Encryption** enables protection of data at rest. A HyperStore system can perform the encryption (and subsequent decryption upon object retrieval) either with a system-generated encryption key (regular SSE) or a customer-provided encryption key (SSE-C).

For more information, visit the Cloudian website:

<https://cloudian.com/>

Lenovo ThinkAgile HX Series powered by Nutanix

Lenovo ThinkAgile HX Series is designed to help simplify IT infrastructure, reduce costs, accelerate time-to-value, and deliver greater reliability. These hyperconverged appliances from Lenovo combine industry-leading hyperconvergence software from Nutanix with Lenovo ThinkSystem™ servers.




Figure 2 ThinkAgile HX3320 (top), HX3520-G (middle), and HX3720 (bottom) appliances

HX Series is built on Lenovo ThinkSystem servers, #1 in both reliability and customer satisfaction. These servers feature enterprise-class reliability, availability, management, security, and performance features. The Lenovo ThinkAgile HX Series enables a feature-rich hyperconverged infrastructure to modernize enterprise data centers.

The HX series consolidates servers, storage, and virtualization into one virtual structure, and easily managed as clusters through a single interface. Storage and compute are pooled across each node in the cluster as a single shared virtual resource for increased and efficient utilization. HX Series simplifies the data center, reduces costs, accelerates time-to-value, and delivers greater reliability.

HX appliances are offered in five workload-optimized series to meet your specific business needs:

- ▶ 1000 Series is optimized for remote and back offices
- ▶ 2000 Series for small to medium businesses
- ▶ 3000 Series is optimized for compute-heavy applications
- ▶ 5000 Series is optimized for storage-heavy workloads
- ▶ 7000 Series is optimized for high-performance environments.

For more information, visit the Lenovo HX Series product page:

<http://www3.lenovo.com/us/en/data-center/converged-systems/Lenovo-Converged-HX-Series/p/WMD00000036>

Nutanix

Nutanix's industry-leading software is integrated in Lenovo ThinkAgile HX series appliances and brings the benefits of hyperconverged technologies to enterprise applications through enterprise storage, data protection, infrastructure resilience, management and analytics, and security.




Figure 3 HX Series architecture

Nutanix simplifies the data center by combining server, storage, networking and virtualization resources into a turnkey hyperconverged appliance. Pre-installed on Lenovo ThinkSystem servers, the HX Series can run any application at any scale and deployed in less than 60 minutes. Users can run large bare metal enterprise applications, like SAP and SQL, alongside VDI and server virtualization apps, like VMware and KVM, simply and efficiently, without compromising performance or manageability.

Nutanix leverages a hybrid delivery model to capitalize on public cloud advantages while retaining the security and control of private data centers.

For more information, visit the Nutanix website:

<https://www.nutanix.com/>

Rubrik Cloud Data Management

Rubrik Cloud Data Management (CDM) platform unifies backup, instant recovery, replication, global indexed search, archival, compliance, and copy data management into a single scale-out fabric across the data center and public cloud. Enterprise organizations benefit from Rubrik because it securely manages all data, physical or virtual, across all locations – on-premises, edge of the data center, and cloud.

A key feature for Rubrik CDM is instant search that delivers near-zero RTOs with predictive search. Easily locate virtual machines (VM)s, databases, applications, or files regardless of whether they reside in the cloud or on premises. The policy-driven engine and programmatic interface eliminate daily operational management by automating how data services are created, consumed, and retired from across the data center and to the cloud. Orchestration is the core of Rubrik's capabilities. It provides a suite of APIs that enterprises can use to manage data, including moving data from the data center to the cloud. Rubrik provides the flexibility for users to provision data management services using configuration management tools or custom portals.

Data is secure in transit and at rest throughout the entire lifecycle, regardless of location. Granular role-based access can be leveraged while automating compliance reporting in order to successfully meet and complete various industry audits.

Analytics and reporting are provided by Rubrik Envision, which unlocks actionable insight across the datacenter and cloud with customizable reports. Users also leverage platform analytics that detail operational efficiency, compliance, and capacity utilization across their infrastructure.

For more information, visit the Rubrik website:

<https://www.rubrik.com/>

Solution Overview

Managing data growth, outside of the primary storage repository, has become increasingly complex when using legacy architectures. Data centers that use traditional infrastructures are difficult to scale. To address rapid data growth, simple, affordable and scalable approaches to data lifecycle and storage management is an imperative.




Figure 4 Enterprise Data Protection with Hyperconverged infrastructure and Object Storage

To tackle this issue, Lenovo has created a software-defined reference solution that is cost-effective, flexible, and easy to manage with limitless scalability. This solution will address current data management needs as well as those needed for data growth. This solution provides backup and long-term data retention for public and private cloud environments.

Figure 4 on page 8 shows:

- ▶ Lenovo Storage DX8200C powered by Cloudian as the highly-scalable on-premise storage target for data protection with support for policy-based migration to the cloud of choice for archive and disaster recovery.
- ▶ Rubrik Cloud Data Management appliance provides end-to-end data management, allows users to securely access data instantly, automates protection policies, and orchestrates data across multi-cloud environments.
- ▶ The Lenovo ThinkAgile HX series powered by Nutanix cluster allows for applications to run in an efficient yet powerful hyperconverged environment where data is initially created and stored.
- ▶ No additional software installation or plugins are needed to connect Rubrik with the Lenovo DX8200C powered by Cloudian Or Lenovo ThinkAgile HX Series appliances

Operational costs associated with backup and recovery are further reduced by combining Rubrik's search and instant mount capabilities with Cloudian HyperStore storage policies that permit users to store and access data in geographically separated regions or data centers. No time is wasted restoring indexes, combing through catalogues, recalling tapes, or even waiting curbside for the tape truck to pull up.

Joint solution capabilities

The Lenovo reference solution provides the following:

- ▶ Automated protection and restore workflows
- ▶ Policy-driven protection and retention operations
- ▶ VM-granular backup and restore
- ▶ Auto-protect newly discovered VMs
- ▶ Export and instantly recover VMs
- ▶ File browse and download
- ▶ Global search, erasure coding, and reporting
- ▶ Secure replication or archival to multiple sites
- ▶ End-to-End encryption and immutability to fight ransomware
- ▶ Start small and scale as you need (Capacity and Performance)
- ▶ Integrated Disaster Recovery
- ▶ This solution supplements disaster recover (DR) solutions to replicate data between on-premises sites, and on-premises to public cloud.

Joint solution benefits

Solution benefits include:

- ▶ Lower TCO with consolidated servers, virtualization, and storage
- ▶ Scalable infrastructure that can scale out as needed
- ▶ Reduction of archival costs to pennies per GB each month
- ▶ Optimized and efficient multi-site data archival and management
- ▶ Ability to orchestrate recovery from disaster via APIs
- ▶ Configurable levels of protection for always available archived data
- ▶ Greater reliability with Lenovo ThinkSystem server platforms
- ▶ Simple manageability through Lenovo XClarity™ management software

The Lenovo reference solution dramatically simplifies how enterprises and service providers manage backups and long-term data retention in hybrid cloud environments.




Figure 5 Solution extensibility for primary storage and scale out

Rubrik CDM delivers a secure and intelligent on-ramp to the Lenovo DX8200C powered by Cloudian object storage platform managed from behind the data center firewall. Rubrik eliminates the complexity of legacy enterprise backup and disaster recovery. It collapses backup software, catalog management, replication, and deduplicated storage into a single set of software that scales linearly. Rubrik's distributed architecture allows deduplication and other data services to linearly scale to maximize efficiency and savings. As shown in Image E, Rubrik sends deduplicated data to the DX8200C storage cluster, reducing data transfer and storage costs.

By combining Rubrik and Cloudian HyperStore storage policies, users can archive data in geographically separated regions or data centers. Cloudian HyperStore is designed for simplicity and durability for easy day-to-day operations and provides the capability of assigning different storage policies granularly, on a per bucket basis. For data protection, users have the flexibility to store data using either replication or erasure coding, according to their needs. Replication results in greater performance, but less storage capacity. Erasure coding provides the most cost-effective way to store data because the data is stored in smaller chunks across the cluster. HyperStore also supports multiple compression algorithms that compress objects on the storage side, saving space and cost.

Enterprises are also deploying Lenovo DX8200C Cloudian storage as an additional cloud storage resource tier in addition to secondary offsite locations for disaster recovery. The solution delivers near-zero recovery times without rehydration to primary systems and global file search for instant data access irrespective of the location.

The solution is designed to be simple and easy to setup and operate, with no additional software or plugins necessary.

To set up the solution, users simply:

1. Configure DX8200C as an archival location, Figure 6.

Add Archival Location

Archival Type

☐ Amazon S3
 ☐ Azure
 ☒ Object Store
 ☐ NFS
 ☐ Tape

Object Store Vendor

S3 Compatible (Basho, Cleversafe, Cloudian, IIJ) ▼

Access Key

Secret Key

Host Name

Bucket Prefix

Number of Buckets

Archival Location Name

S3Compatible-Cloudian

RSA Key ⓘ

Cancel Add

Figure 6 Archive location selection in Rubrik interface

2. Add DX8200C as an archival target by entering the Access Key and Secret Key for the DX8200C user, and entering the RSA key and hostname.
3. Once DX8200C is configured and available to select as an archival location, create or edit an SLA Domain and define the archival location, Figure 7 on page 12.

Create SLA Domain

Remote Storage Configuration

Retention On Brik

30 days 2 years

Archival Location

S3-CloudOn-NorCal

S3COMPATIBLE-CLOUDIAN

☐ Enable Instant Archive ⓘ

Archival starts after data is expired locally, and is retained for 1 year 335 days.

Replication

A replication target has not been set up yet. Please [add a replication target](#) to configure retention.

Cancel Configure SLA Create

Figure 7 Sample Rubrik UI of SLA Domain configuration

4. Move the slider to specify when data should be archived to the DX8200C. By default, the retention period on Rubrik is 30 days but is easily be customizable to meet the business or governance requirements.

Once DX8200C has been added as an archival location and the SLA Domain has been configured to archive to that location, the following information may be easily found in the Rubrik user interface, Figure 8 on page 13

- ▶ Amount of data archived
- ▶ Amount of data downloaded
- ▶ Number of objects archived




Figure 8 Sample view of Data Management

Solution testing and validation

This solution was tested extensively in the Lenovo Innovation Center.

The HX Series appliances hosted the software that simulated business application workloads that generated the data to be protected, backed up and recovered. The DX8200C appliances provided the enterprise-class object storage used to store and archive the data. Rubrik software provided the policy-based management for the automated backup and instant recovery of the data to/from the HX Series and DX8200C appliances.

The objectives for testing are as follows:

- ▶ Validate solution manageability,
- ▶ Confirm end-to-end operation and functional capabilities.
- ▶ Document benefits including but not limited to the fundamental protection of the application data and also the restoration of that data from a synthesized failure situation.

A typical configuration was developed and included the following solution elements:

- ▶ Four node HX clusters: HX Series appliances powered by Nutanix
- ▶ Four node DX8200C cluster: DX8200C appliances powered by Cloudian
- ▶ Three node Rubrik Cloud Data Management appliance (Rubrik r334 appliance)
- ▶ Two Lenovo Top-of-Rack switches: RackSwitch™ G8124E and G8272

The bill of materials the hardware is listed in Table 1.

Table 1 Configuration details

Component	Quantity
DX8200C node hardware configuration details	
Lenovo Storage DX8200C 5120AC1	1
Cloudian HyperStore SW stack for DX8200C	1
Processor: Intel Xeon Processor E5-2630 v4 10C 2.2GHz 25MB Cache 2133MHz 85W	1
Memory: 16GB TruDDR4™ Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	4
3.5 Inch Hard Disk Drives (max. 14): 4TB 7.2K 6Gbps NL SATA 3.5" G2HS HDD	14
SFF Slim SAS SATA (max. 2): Intel S3510 480GB Enterprise Entry SATA G3HS 2.5" SSD	2
PCI-E x8: Intel x520 Dual Port 10GbE SFP+ Adapter	1
HX node hardware configuration details	
HX3510-G Nutanix Appliance	1
Nutanix SW Stack on Nutanix Acropolis	1
Intel Xeon Processor E5-2680 v4 14C 2.4GHz 35MB Cache 2400MHz 120W	2
32GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	8
1TB 7.2K 6Gbps NL SATA 2.5" G3HS HDD	6
Intel S3610 800GB Enterprise Mainstream SATA G3HS 2.5" SSD	2
Intel x520 Dual Port 10GbE SFP+ Adapter	1
Switch details	
Lenovo RackSwitch G8124E, 24 port, 10Gigabit	1
Lenovo RackSwitch G8272, 48 port 10GbE & 6 port 40GbE	1

Use Cases

This section is intended to provide a few deployment examples of the solution. It is not intended to be an exhaustive list, merely a set of sample use cases.

Archival

Over the past several decades, corporate data has grown at a record pace. Today's backup and archive targets — tape, disk, and cloud — all have tradeoffs to consider in using them to keep up with the growth, cost-effectively and reliably.

Resulting issues of these tradeoffs include excessive backup times, especially when backups fail, and the challenge of meeting service level agreements (SLAs), along with the Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO).

Lenovo DX8200C users store unstructured data across one or multiple buckets depending on storage policies. Rubrik provides the ability to have multiple policies which can write data across multiple buckets - making sure the data is archived to the most appropriate location ensuring success on the objectives that matter most: backup window predictability and repeatable RTO/RPO.

Building a 3-year cost schedule makes savings even more apparent across different types of storage media. For example, protecting 100TB of source data with an archival policy comprised of 30 daily, 12 monthly, and 3 annual snapshots, users realize 1.7x to 2.0x hard savings when comparing on-site DX8200C vs. NAS backup target or tape archival. See Figure 9.




Figure 9 Data Archival Use Case

The following assumptions have been made:

- ▶ Rubrik is used to archive to DX8200C
- ▶ Rubrik is used to archive to NAS backup target
- ▶ User has opted for LTO-5 tapes with daily pickup

- Monthly cost of \$.01 per GB for Clodian HyperStore, \$.60 per GB for NAS backup target, and \$30 per LTO-5 tape

The solution not only simplifies how enterprises and service providers manage backup and long-term data retention in hybrid cloud environments but also dramatically reduces TCO associated with long-term data archival.

Integrated Disaster Recovery

In a Disaster Recovery (DR) scenario, users leverage a variety of options often combining various forms of replication, DR automation tools, and backups. This can be costly because it requires similar hardware at production and the DR location site, as well as the complexity of managing multiple tools.

A good DR plan provides multiple layers of protection by leveraging tools that protect from a myriad of failure types. This can be accomplished with tools such as snapshots, data replication, and data protection. Rubrik plays a key role in DR planning by providing one policy engine to setup backup, replication, and archival schedules on-site or in the cloud. Rubrik enables hourly RPO and near-zero RTO.

By leveraging components of a modern data center such as a converged architecture, flash, and deduplication, data can be efficiently replicated from Rubrik instance to another with VMs instantly recovered.

Lenovo DX8200C powered by Clodian also offers a distributed, peer-to-peer architecture with no single point of failure. Having a system that is resilient to network and node failures without data loss greatly enhances DR capabilities.

The DX8200C also offers a unique way of making DR more resilient and cost efficient with options ranging from replication of data across nodes or data centers, or utilizing the distributed erasure coding configuration which optimizes and protects the data across single or multiple data centers. For users with one to two data centers, replication using RF3 is recommended. For larger organizations with more than two data centers, distributed erasure coding configuration is recommended within Clodian HyperStore.

Scenario 1 - Public cloud

Disaster recovery is complicated. With the increasing adoption of public cloud, enterprises are looking for ways in which the use of compute on cloud can be innovatively leveraged. One such way is to leverage the cloud for DR, instantiating on-premises VMs to AWS EC2. The solution, combining Rubrik and the Lenovo DX8200C with Clodian makes it easier to failover to the cloud. Instances in the cloud can be spun up on demand, as required.




Figure 10 Disaster recovery with public cloud use case

Rubrik's CloudOn feature leverages data that has been tiered to cloud through Cloudian HyperStore and automates the process of converting on-premises virtual machines into cloud instances running your applications in the cloud. The ability to power on a workload in the cloud is available for backups taken of VMware vSphere VMs that have been tiered to Amazon through Cloudian. There is no requirement to have a Rubrik Cloud Cluster running in the target Amazon region.

Using the Lenovo DX8200C powered by Cloudian's flexible Amazon S3 compatible object store, Rubrik can instantiate the snapshots that are archived to S3 efficiently. If snapshots are archived to Cloudian, Rubrik will first bring the snapshot on-premises and then move it to S3 for instantiation.

Rubrik offers three options that can be applied to a workload being archived from on-premises to Amazon S3:

- ▶ On-Demand. The default configuration in which AMIs are created only at the time of a "power on in the cloud" request.
- ▶ Auto Convert Latest Snapshot — Keep One. Rubrik will automatically construct an AMI reflecting the latest snapshot to be archived into S3. When a new snapshot is sent to the archive, a new AMI is constructed with the new archive data. Once completed, the older AMI is removed.
- ▶ Auto Convert Latest Snapshot — Keep All. Rubrik will automatically construct an AMI reflecting the latest snapshot to be archived into S3. When a new snapshot is sent to the archive, a new AMI is constructed with the new archive data. The older AMI is retained if desired (configurable via policy) creating a series of AMIs representing each snapshot.

These options provide flexibility for ensuring RPO + RTOs are met. Continuous conversion of snapshots can be used to enable smaller RPO + RTO windows. The option to spin up instances on demand as required, can enable business to reduce costs by avoiding the need for a DR data center and dedicated DR infrastructure.

Scenario 2 - offsite replication

In this age of extensive data expansion, IT organizations are finding it increasingly challenging to:

- ▶ Easily scale storage and implement DR as demand increases.
- ▶ Ensure 100% data availability across multiple data centers.
- ▶ Provide seamless data access in the event of a node or even data center failure.




Figure 11 Disaster Recovery site/offsite use case

The solution deployment models ensure efficient data backup and recovery options, ranging from replication across nodes within a single data center, replication across multiple data centers within a single region, and distributed erasure coding across multiple data centers and multiple regions. Lenovo DX8200C with Cloudian HyperStore may be used to expose multi-site configurations and S3 endpoints written to by Rubrik.

Archives stored on Cloudian HyperStore can withstand a loss of a single data center while still ensuring data protection and availability. This is done by:

- ▶ Replicated erasure coding across regions and data centers.
- ▶ Distributed erasure coding solution across different regions and data centers.
- ▶ Data accessibility via multiple endpoints enabled by Cloudian HyperStore's unique implementation of erasure coding (ranging from EC2+1 to EC12+4) and configurable durability.

The Lenovo DX8200C powered by Cloudian's replicated and erasure coding configurations with Rubrik's Cloud Data Management, can be augmented to ensure constant data availability.

Scenario 3 - Offsite for natural disasters

Power outages, hurricanes, and other forms of natural disaster happen. But what happens if an entire datacenter is lost? Legacy backup vendors do not store metadata with archive, making restoring from archive nearly impossible.

It is costly for a business to maintain a DR site that mirrors the production site's full capabilities. Commonly DR sites are a stripped-down version of the production site, therefore

a user may not maintain Rubrik in both sites. The good news is that Rubrik stores backup metadata along with the archive. This means that, should the primary site be unavailable, full restores can be easily accomplished from the Cloudian archives once a Rubrik physical appliance or virtual Edge appliance is connected.

The security of data is prioritized and therefore stored in an immutable format readable only by Rubrik. This dramatically increases the ability to recover from security breaches such as ransomware. Of course, with the data format readable only by Rubrik, Rubrik software must be in place to recover the archived data from the DX8200C archive.

Conclusion

The integration of Lenovo, Cloudian, Rubrik and Nutanix capabilities create a robust solution that addresses

- ▶ Data center simplification
- ▶ Cost-efficiency
- ▶ Data backup, archive and disaster recovery.

The solution delivers a modern infrastructure with limitless scalability, security, and flexibility to expand as needed over time. Data is efficiently managed either on-premises, in the cloud, or in a hybrid model.

Combining the Lenovo ThinkAgile HX Series powered by Nutanix hyperconverged appliances, along with Lenovo DX8200C appliances powered by Cloudian with Rubrik enables a highly scalable and fault tolerant backup and archive storage and a robust application infrastructure. This solution helps to simplify how enterprises and service providers manage backup and long-term data retention in hybrid cloud environments, while enabling seamless scalability for the future.

Authors

Grant Jacobson is Director of Alliance Marketing at Cloudian. He has 20 years of experience in IT, spanning enterprise software, data storage, compute and networking. He holds Bachelors and Masters degrees from the University of Texas.

Bill Hansen is a Senior Product Manager and Product Planner at Lenovo. His current focus is on software defined storage platforms. He has over 15 years of experience in the storage industry including roles in storage engineering and product management.

Thanks to the following people for their contributions to this project:

- ▶ Rachel Zhu, Lenovo
- ▶ Anu Nair, Cloudian
- ▶ Rebecca Fitzhugh, Rubrik
- ▶ Demetrius Malbrough, Rubrik

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on January 9, 2018.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/lp0821>

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	Lenovo(logo)®	TruDDR4™
Lenovo XClarity™	ThinkAgile™	
RackSwitch™	ThinkSystem™	

The following terms are trademarks of other companies:

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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