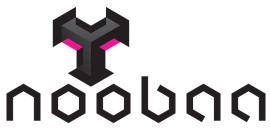

NooBaa and Lenovo Reference Architecture

Last update: May 2018

Provides a highly scalable data platform system

Eran Tamir [NooBaa]
Yuval Dimink [NooBaa]

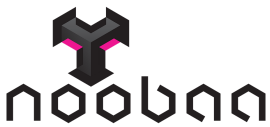

Table of Contents

1 Table of Contents

2 Introduction

This document describes a reference architecture for NooBaa. It was created to assist in planning, designing, and implementing the NooBaa software on Lenovo hardware.

2.1 Target Audience

The intended audience of this document is IT professionals, solution architects, sales engineers, DevOps and partners. This paper assumes that the reader has administrator level knowledge of IT infrastructure (servers, network storage) and some user level knowledge of cloud services.

3 Executive Summary

Massive growth of unstructured data in almost every industry demands a flexible, easy to scale, and simple solution to manage such. The cloud introduces a great opportunity to lower cost and easily protect data by moving it offsite but it does not integrate easily into on-premises solutions or between cloud providers.

NooBaa addresses those needs by allowing enterprises and MSPs to store unstructured data in a simple, flexible and scalable manner. Data is accessed using standard Object Storage cloud API's. NooBaa's data platform can use storage resources from any combination of x86 servers, existing storage products and natively integrated cloud services such as AWS S3, Azure Blob and GCP storage. This grants a more economical and manageable storage fabric compared to existing products.

NooBaa fabric is elastic and can easily be scaled in any increment and in any direction or location and administrators have the ability to store data across multiple data centers or geographically separated regions.

NooBaa is a certified backup target for applications like Commvault, NetBackup, and more.

4 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.

4.1 Business problem and business value

Unstructured data grows at a faster pace each year. Organizations are required to accommodate to how this data is generated and consumed as well as to how it is stored and managed. In many cases shadow IT, M&As and changes in strategy create a reality of on-premises and in the cloud storage silos. The cloud introduces additional challenges in terms of security, privacy and data gravity.

These require a data storage solution that can serve any workload, anywhere as needs arise and the cloud must be natively integrated into the solution in a transparent manner. Moreover, the ability to change easily, due to regulatory or business requirement, becomes mandatory and thus the solution must be elastic and easy to manage.

4.2 Business value

To solve the massive growth of data storage, multiple silos on-premises and in the cloud and the regulation challenges, organizations are looking for new approaches to data management. Such solutions typically have the following key features:

- Scale out architecture
- Hardware agnostic
- Hybrid and multi-cloud data management
- Standard cloud services API, such as AWS S3 and Azure Blob compatible API for Object Storage solutions
- Multiple data centers, cloud multi region, etc.
- Data efficiency such as compression, deduplication, erasure coding.
- Data management policies

5 Use cases

This solution fits well in many use cases, and is particularly appealing for the following use cases:

5.1 Active Active Archive

Active archiving lets customers keep the data available for their applications, while the data can be anywhere.

This use case is mainly relevant for workloads that produce a lot of data on-premises, usually on fast and expensive storage, while a large portion of it will never be used.
In this case, the data will be stored on NooBaa, while a lifecycle policy or a custom data flow function will move local copies to the cloud. The **data will be available regardless of the actual location**, and NooBaa will serve it, with a locality preference.

5.2 Video surveillance

Video surveillance is a classic example of data that is not accessed on a regular base but when the need arises, the data is needed in real-time. In this use case, NooBaa will perform as the primary storage for video catalog. The **deduplication algorithms** will keep the data storage efficient, mainly for the periods where there is no motion.
The active archiving can take place, **move old videos to the cloud, transparently**, so any application that tries to read data will get it, regardless of location. A metadata extraction to feed any database with important characteristics extracted from a movie can take place on the fly. When a certain video is required, no restore process is needed as NooBaa will serve the data regardless of the location.
NooBaa's data resiliency helps NooBaa's customer keep the budget low by maintaining **multiple data resiliency policies** in a single repository. Every class of data can have a different resiliency such as different level of erasure coding

5.3 Media and Entertainment

NooBaa can connect directly to any Media Asset Management (MAM) system that works with AWS S3 API and serve the media catalog seamlessly. The data storage can be on-premises, easily replicate to multiple locations for the various production phases, and later on, transferred for distribution. The entire data flow can be automatic, based on dates, extensions or any other relevant characteristic including custom metadata.
Extraction of metadata information to feedback the MAM system with important

characteristics can take place on the fly. In the same way, NooBaa can invoke any API of tools that are used for face recognition, subtitles creations, and much more as part of this custom data flow.

5.4 Backup Target

NooBaa supports both AWS S3 compatible API and Azure Blob compatible API. Any backup or archiving software that can use these API will work with NooBaa seamlessly.
Using NooBaa in this way let customers use any mixture of on-premises, hybrid cloud storage or public cloud-native storage. In addition, customers can easily set a unique data placement policy for every application, or the data you want to backup or archive.

NooBaa has official certification from Veritas and Commvault, but also successfully tested with Rubrik, CloudBerry, Synology, Cyberduck and more.

5.5 Storage as a Service

Managed Service Providers can offer Storage as a Service by using NooBaa technology. With NooBaa, MSP can turn existing data centers into a Storage-as-a-Service and start providing ideal and cost-effective backup and DR services with S3 compatibility.

Using AWS S3 API, the de facto standard, allows you to utilize the cloud ecosystem including backup applications, archival solutions, and a documented API. With its true hardware agnostic technology, NooBaa lets you aggregate multiple storage silos of any size and vendor. Due to NooBaa's unique architecture, MSP can scale the storage quickly and anywhere.

NooBaa's management has a rich API to enable a quick and smooth integration with billing, account creation, permissions etc.

5.6 Cloud Native Data Storage

NooBaa enables its customers to easily embrace cloud flexibility and agility within their data centers.

The multiple storage silos that exist in most data centers today make it complex and costly to manage. By nature, each silo enforces its own lifecycle, scale, hardware refresh and limitations. NooBaa introduces a different experience, turning

existing storage infrastructure into a private cloud with built-in agility and flexibility, with software only. It reduces many risks in terms of resource investment and eliminates the limitations on scaling and geographic data distribution and the pain of migration and hardware refresh.

With AWS S3 compatibility and hardware-agnostic technology, NooBaa lets customers easily aggregate multiple storage silos of any size and vendor.

5.7 IoT

IoT has by nature a distributed architecture, where multiple devices send bits to TBs of data to centralized locations on a daily basis.

With NooBaa's technology, customers can build multiple data centers, close to clusters of devices and provide the primary storage that handles the data aggregation. Using NooBaa's independent serverless functions, data manipulation, aggregation, de-identification, and much more can take place automatically and then transfer the data for analytics and diagnostics. The classic example is autonomous cars that collect data from multiple sources.

NooBaa can be used in any country, collect all the data, de-identify it, per GDPR regulation requirements, compress, encrypt and transfer the data to a centralized location for analytics.

5.8 Medical imaging

Medical imaging introduces regulatory and data gravity challenges.
NooBaa's technology helps customers to create an on-premises storage that can store the medical imaging for the diagnostic phase. Once the data is cold, a de-identification process can automatically mask the DICOM records, update a database with the connecting link and move the local copies of the encrypted data to the cloud using NooBaa's serverless functions. Such a customized data flow is just one example of personalized data flow that any customer can create to better tailor NooBaa to its need.

5.9 Research projects

Research projects are always tricky. The amount of raw data is always a challenge while trying to analysis it. In many cases, the same huge dataset is used for multiple projects, while on-premises computational resources are limited. NooBaa helps customers to automatically split the relevant chunk of data out of the raw

data and easily move it to the cloud, any cloud, for analysis. This method reduces the amount of data you need to push to the cloud, de-identifies or masks the data if needed, and effectively uses the right computational resources for the project lifetime.

The method helps to boost research, keeps privacy in place, squeezes timeframes, and stays in budget while using only temporary resources instead of investing in computational resources that may never be used again.

6 Supported Requirements

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

6.1 Supported Functional requirements

Table 1 lists the functional requirements for a cloud object store.

Table 1: Functional Requirements

Requirement	**Description**
Multi-tenancy	Accounts and Roles.
Data placement policy	Store data in geographically dispersed regions or native private/public cloud storage.
Billing and reporting	API to read all usage and quota information per bucket or account.
Encryption	AES 256 encryption of all objects, [Optional] client-side encryption with AWS SDK, based on a master key.
Monitoring	Syslog and detailed exportable audit.
Data Resiliency	Choose different protection levels, such as replication or erasure code.

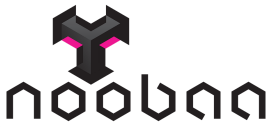

6.2 Supported Non-functional requirements

Table 2 lists the non-functional requirements for hybrid and multi-cloud data platform

Table 2: Non-functional Requirements

Requirement	Description
Scalability	Solution components scale for growth without losing performance.
Data resiliency	Self-healing of the data, erasure coding or full copies of the data.
Physical footprint	Slim solution.
Ease of installation	Quick deployment, stacking and cabling.
Ease of management	Intuitive management to maintain PB scale.
Flexibility	Store the data where it's needed – single or multiple locations, on-premises or in the cloud. No Cloud vendor lock-in, no HW vendor lock-in, can scale with any HW vendor in any increment and benefit from the heterogeneity of the solution.
Security	Encrypt all data at rest and in motion.
Data Efficiency	Deduplication and compression.
Power Efficiency	Uses low-power CPU's and power efficient drives with large capacities.
Cloud integration	Data can be mirrored seamlessly to AWS, Azure, Google or any AWS S3 compatible storage like Wasabi, CloudItalia,etc. NooBaa will serve the data from any available resource, without any awareness of the application.

7 Architectural overview

This section presents an architectural overview of the NooBaa solution and compatibility with the Amazon S3 API.

7.1 AWS S3 API compatibility

AWS S3 is the de facto standard of object storage, hence it is the go-to choice for most modern applications.

The NooBaa storage platform is AWS S3 compatible and allows developers to use

AWS S3 SDK. Having the ability to use the same S3 API with NooBaa across vendors eliminates the need to rewrite applications when backend vendors change.

NooBaa can also automatically tier data between on-premises and cloud deployments while representing the entire data set in a single namespace. With this advanced functionality, NooBaa provides a compatible data platform for on-premises and hybrid cloud deployments.

7.2 NooBaa architecture

NooBaa® software uses a distributed architecture with no single point of failure. It easily scales and shrinks horizontally, vertically and into the cloud by using any hardware, so deployments can start with any starting point such as few servers in a single data center and then scale out as usage increases to thousands of servers distributed across multiple data centers and managing petabytes of data. As another example you can start with pure cloud setup and combine data center resources at later stages. NooBaa's distributed architecture is designed to provide a trusted resiliency by using self-healing capabilities as well as flexible data mirroring that can take advantage of public or private cloud storage. **Figure 1** shows an overview of NooBaa that uses Lenovo servers.

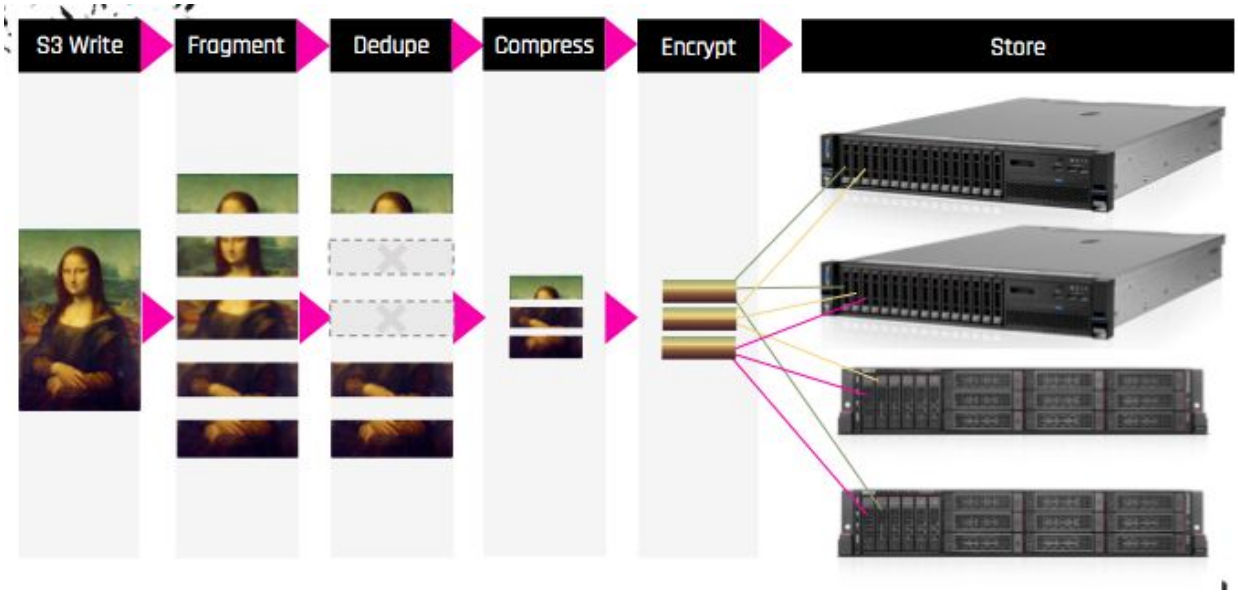

Figure 1 NooBaa and Lenovo architecture

NooBaa dedupes the data using Rabin Sliding Window Dedupe algorithm.

NooBaa compresses objects on the storage side to save space while maintaining streaming performance by utilizing Google's Snappy compression algorithm.

NooBaa encrypts the data using AES 256 encryption.

NooBaa's data resiliency supports both replications as well as erasure coding data, where the tradeoff is between storage efficiency and read/write/rebuild performance. NooBaa provides supporting information to help the admin chose the right policy per class of data. Data resiliency policy that is configured per bucket and can be adjusted over time. Data is self-healed in case of failure and an administrative alert/corrective action is provided.

7.3 NooBaa's Component model

NooBaa is comprised of several components as shown in Figure 2. This section describes the functionality of each component as well as its role as part of the overall architecture.

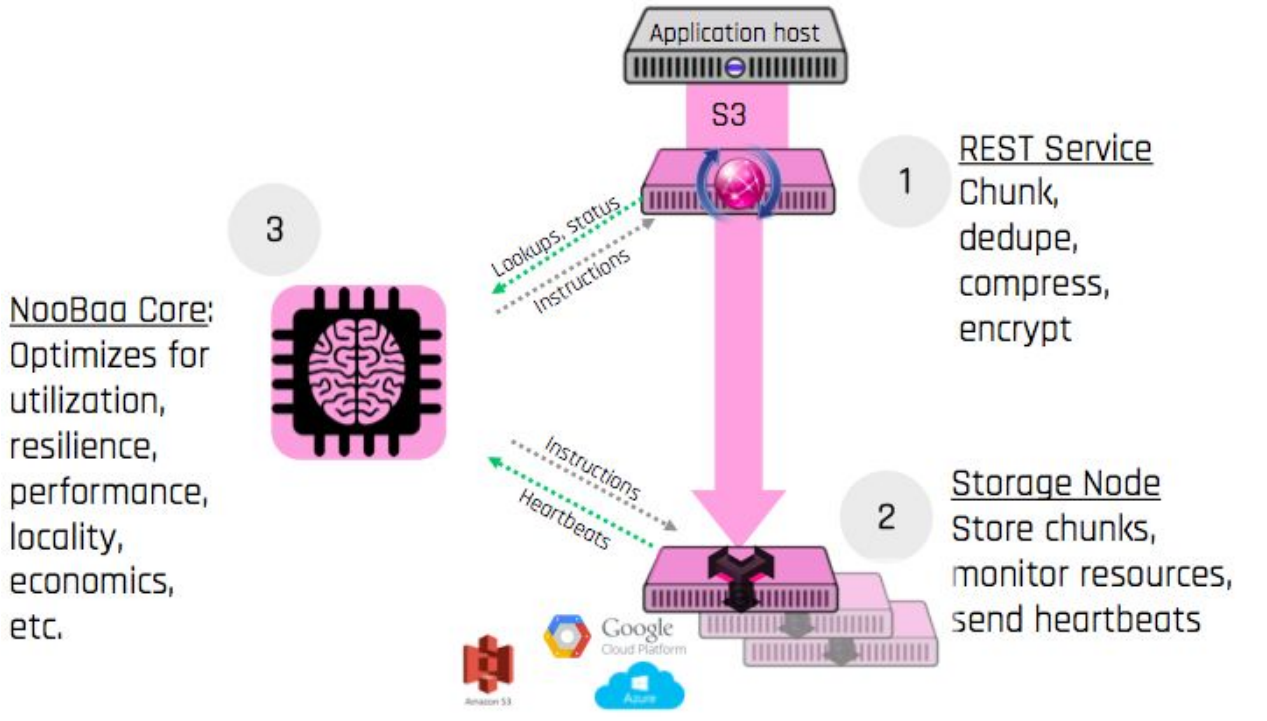

Figure 2 The NooBaa component model

A NooBaa system is comprised of the following components:

7.3.1 NooBaa Core [Required component]

The NooBaa core is packaged as a complete VM appliance that can be hosted in any hypervisor (See shape 3 in Figure 2 The NooBaa component model). As soon as it is deployed and configured, the system is fully functional, and can erve read and write operations, without any additional storage nodes.

The NooBaa core component is a required component for any NooBaa implementation and is the starting point for any such deployment. The NooBaa core is critical to the health of the system, and multiple cores can be configured to form a highly available cluster for resiliency. Scaling from a single NooBaa core to a clustered NooBaa cores does not change the system's management as it is done by the management console across the cluster.

The NooBaa core monitors the entire system, it's services, storage nodes, storage drives, network availability, core cluster as well as the stored data. It allocates storage for write operations using machine learning algorithms that optimize performance, location and cost. It redirects to relevant storage nodes on reading operations.

In production deployments, the data does not flow through the NooBaa core, and it is designed for scale and high availability.

The NooBaa core hosts multiple internal services such as the NooBaa metadata and monitoring services and clustering services.

In addition, it hosts the externally available default Rest Service as well as NooBaa's Web Management Console.

7.3.2 NooBaa's Web Management Console [hosted component]

NooBaa's web management console is the user interface for NooBaa's system administrators, It is hosted within the NooBaa core and is available as soon as a NooBaa core was deployed.

Using the web management console, administrators can configure and manage every function of the system including resources, data buckets, users and more. In

addition to that, using the management console the admin can access and generate various audit

7.3.3 NooBaa Daemon [Optional component]:

The NooBaa daemon is a userspace application that can be installed on any Linux or Windows server. The NooBaa deamon monitors the system that it is installed in and send periodic heartbeats to the NooBaa core with rich monitoring information.

The NooBaa daemon is not required for a functioning system, but a production deployment is expected to have multiple daemons installed, as a function of the number of access points and storage nodes. Installing a daemon is achieved through running one liner command available from the management console.

The NooBaa daemon accepts instructions from the core and communicates with it's peer daemons using p2p communication.

Each daemon is securely signed to be used by a single system.

The daemon hosts two separate services that can be independently turned on and off using the management console:

- The REST services
- The Storage node manager

7.3.4 REST Service [hosted component]

The NooBaa system provides a high-performance object storage service. The REST service serves HTTP R/W requests for the supported API's incoming from authorized clients. The REST service is hosted within a NooBaa Deamon and can be turned on and off.

The service is stateless and scalability is achieved by enabling it on multiple NooBaa daemons.

The REST service also offloads the deduplication, compression and encryption of the data, and writes chunks of processed data to the storage nodes using P2P.

Clients are authorized using a combination of access and secret key.

7.3.5 Storage Node [hosted component]

The Storage Node service is hosted within the NooBaa daemon. The storage node monitors the local filesystems capacity and performance, stores and serves encrypted data chunks using P2P and reports on any data tampering or bitflips that require self healing.

8 Operational model

This section describes the NooBaa operational model that was verified with Lenovo hardware and software. It concludes with some example deployment models.

8.1 Storage servers

NooBaa can be used with any server. The below servers were listed due to their good fit to the platform.

8.1.1 Lenovo ThinkSystem SR650 Server

Lenovo ThinkSystem SR650 (as shown in Figure 3) is an ideal 2-socket 2U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR650 server is designed to handle a wide range of workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), enterprise applications, collaboration/email, and business analytics and big data.

Figure 3 Lenovo ThinkSystem SR650

For more information, see the following websites:

- ThinkSystem SR650 Datasheet
- Lenovo ThinkSystem SR650 Server Product Guide

Server configuration:

Drives (Hot swappable)	12(3.5") x 12TB or 24(2.5") x 7.68TB
Total Storage Capacity	144TB or 184TB
Internal SSD	2 x 480 GB SSD
Data Protection	Raid 1 for internal SSD
Network Interface	10 GbE dual port NIC
Form Factor	2U
CPU	Intel Xeon E5-2620 v4 Series
Memory	64 GB
Monitoring/Management	CLI, GUI, API, IPMI, JMX

8.2 Network switches

NooBaa can be used with any network switch. The below switches were listed due to their good fit to the platform.

8.2.1 Lenovo RackSwitch G8124E

The Lenovo RackSwitch G8124E (as shown in Figure 5) 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 4 Lenovo RackSwitch G8124E

For more information, see the RackSwitch G8124E Product Guide.

8.3 Capacity planning

8.3.1 Storage server capacity planning – general description

The NooBaa system creates internal data redundancy is not dependant on any sort of RAID configuration on the storage servers. NooBaa supports both multiple data copies as well as erasure coding for data redundancy.

Data redundancy policies are configured per data bucket to optimize for efficiency, availability and cost per workload. The UI dashboard provides an easy way to understand available and usable storage.

It's important to keep in mind that there is a trade-off between data capacity efficiency and rebuilt time. For example, 3 copies will provide 200% storage overhead, tolerance for 2 failures, with a quick rebuilt time in case of a failure. Erasure coding of 10+4 would have overhead of 40%, tolerance for 4 failures, while rebuild time will be very high.

Data redundancy is achieved to support disk failure but server failure can also accrue. As that is the case, the system will spread the data between servers, given large enough number of servers to allow it.

Either way, as soon as a failure occurs, the system will identify the failure, will alert about it and will trigger self-healing to bring back the data to full redundancy within the available resources.

1.1.1 Storage server capacity planning – multiple copies examples

For data replication, the recommended number of copies is 3. This means that for the original data copy, two additional and identical copies of the data are created in other drives/servers and thus the overhead in such scenario is 200%.
For example: storing 200TB of data will require 200TB+400TB=600TB of overall storage

Figure 5 table lists the number of servers that are required per each capacity, assuming 3 copies per data chunk (200% overhead as described above)

Drives per server	12	12	12	60	60	60
Drive Capacity	12 TB	12 TB	12 TB	12 TB	12 TB	12 TB
Total server capaciy	144 TB	144 TB	144 TB	720 TB	720 TB	720 TB
Total number of servers	1	2	3	1	10	20
Total capacity	144 TB	288 TB	432 TB	720 TB	7,200 TB	14,400 TB
Usable capacity with 3 copies	48 TB	96 TB	144 TB	240 TB	2,400 TB	4,800 TB
Server failure redundancy	No	Yes	Yes	No	Yes	Yes

Figure 5 NooBaa multiple copies sizing

1.1.1 Storage server capacity planning – Erasure coding examples

For erasure coding, any ratio between data and parity can be configured with a recommendation of 4+2 due to the tradeoff for rebuild time. This means that the the original data set will be encoded with an additional 33% for parity and 2 drives can be lost without affecting data.

For example: storing 200TB of data will require 200TB+100TB=300TB of overall storage

Drives per server	12	12	12	60	60	60
Drive Capacity	12 TB	12 TB	12 TB	12 TB	12 TB	12 TB
Total server capaciy	144 TB	144 TB	144 TB	720 TB	720 TB	720 TB
Total number of servers	1	2	3	1	10	20
Total capacity	144 TB	288 TB	432 TB	720 TB	7,200 TB	14,400 TB
Usable capacity with (4+2)EC	96 TB	192 TB	288 TB	480 TB	4,800 TB	9,600 TB
Server failure redundancy	No	Yes	Yes	No	Yes	Yes

1.2 Networking planning

This section describes the networking topology, including a design guidance 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 data communication services. If more ports are required, the 64-port Lenovo RackSwitch G8264 switch can be used. The management interface on the storage 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 6.

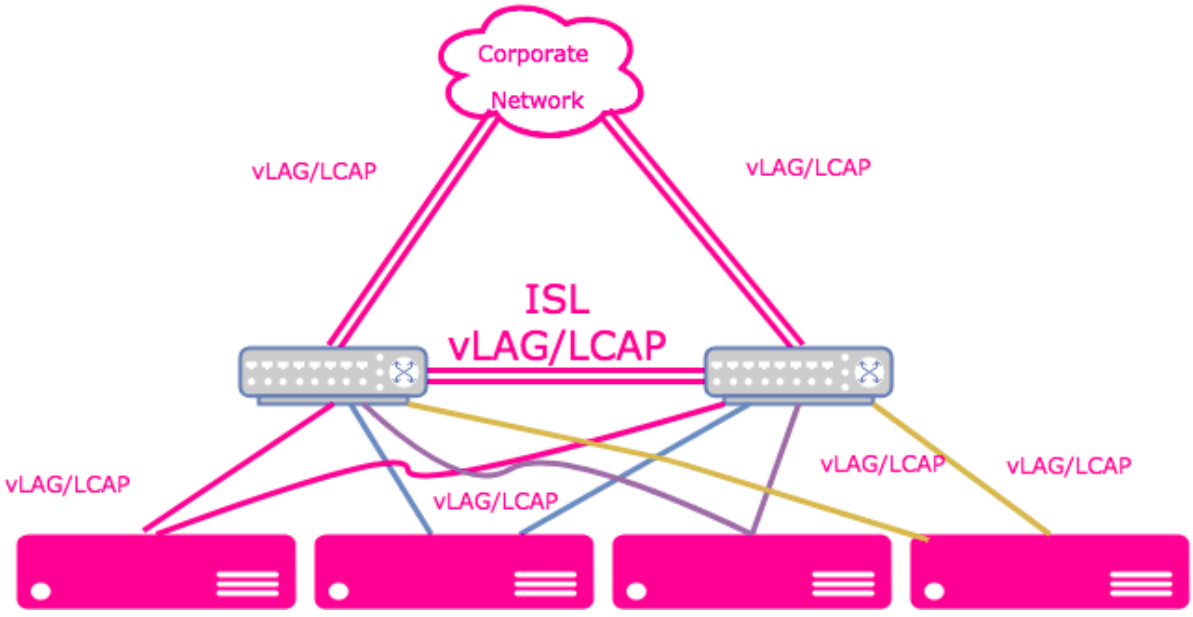

Figure 6. LACP/vLAG recommended network design

An Inter Switch Link (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, also known as a dynamic trunk group. When a link in a LACP trunk group fails, traffic is redirected dynamically to the available 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 NooBaa.

1.3 Data architecture

Each NooBaa software implementation is recommended to start with three or more storage nodes, serving full data replication or erasure coded across the available storage nodes for data durability and availability. The number of replicas or erasure code ratio required to meet SLA and cost objectives, is configurable per bucket, including the option to mirror data to other data centers or the cloud for geo redundancy. Reads and writes operations will get locality priority and can be served by multiple stateless endpoints.

Figure 7 shows an example of a NooBaa cluster that is distributed and elastic

across Geos.

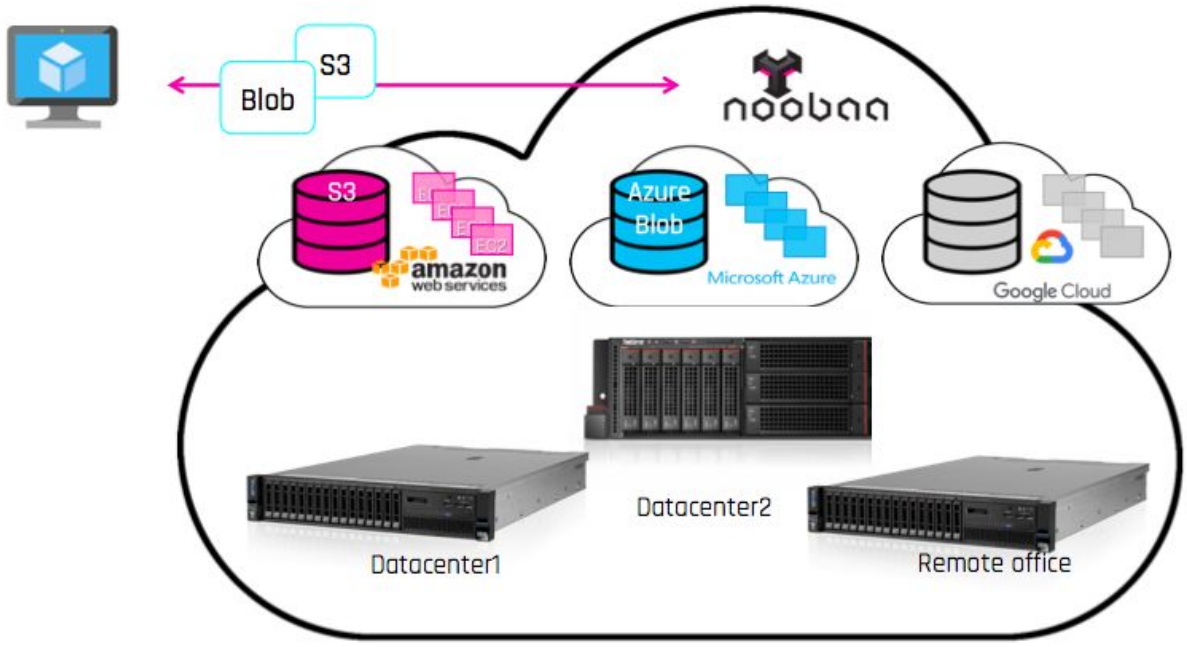

Figure 7 NooBaa data architecture

2 Deployment considerations

This section describes several considerations, such as high availability, load balancers, endpoints, usage of cloud resources, etc.

2.1 High availability

The Lenovo solution with NooBaa contains high availability attributes. By default, every NooBaa brain includes all the services: REST service, Metadata and monitoring service, as well as infrastructure services like NoSql engine.

NooBaa's unique architecture has independent high availability for the data and the metadata.

2.1.1 NooBaa Core

A cluster of odd number of NooBaa Cores will provide high availability for the metadata and monitoring service.

2.1.2 Storage nodes

The number of storage nodes and drives will determine the availability of the data.

2.1.3 Access points

Multiple REST services access points can be installed anywhere independently and provide a highly available service.

2.2 NooBaa Management Console

The NooBaa Management Console is served by NooBaa brain's core service and doesn't need any installation. When the brain installed in a cluster mode for high availability, management requests will be redirected automatically to the master brain.

2.3 Load balancers

The NooBaa software works with most load balancers that are available on the market. The only service that is recommended for load balancing is the REST Service (S3 and Blob).

2.4 Backup

A backup for NooBaa brain's Virtual disk is recommended. Multiple 3rd party applications offer VM backup, replication, etc. and would work seamlessly with NooBaa. There is no need to back up the storage nodes, as they data placement and resiliency policies will take care of that.

2.5 Tiering

NooBaa allows users to move, manipulate, mask objects on a per-bucket level basis to Amazon S3, Azure Blob storage, Google storage, other NooBaa instance ad literally to any storage on-premises or in the cloud.

3 Appendix: Bill of Materials

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

3.1 BOM for storage servers

Bill of materials for a single Lenovo System x3650 M5 storage server.

Part #	Description	Quantity
5462D4x	Lenovo System x3650 M5, Xeon 8C E5-2630v3 85W 2.4GHz/1866MHz/20MB, 1x16GB, 0/Bay HS 3.5in SATA/SAS, SR M5210, 750W p/s, Rack	1
00FK643	Intel Xeon Processor E5-2630 v3 8C 2.4GHz 20MB Cache 1866MHz 85W	1
46W0796	16 GB TruDDR4 Memory (2Rx4, 1.2V) PC3-17000 CL152133MHz LP RDIMM	3
00FN173	6TB 7.2K 6Gbps NL SATA 3.5in G2HS 512e HDD	12
00FK658	System x3650 M5 Rear 2x 2.5in HDD Kit	1
46C9114	ServeRAID M1215 SAS/SATA Controller	1
00YC325	S3710 400GB Enterprise Performance SATA G3HS 2.5in SSD	2
00KA498	System x3650 M5 PCIe Riser (2 x8 FH/FL + 1 x8 FH/HL Slots)	1
00JY820	Emulex VFA5 2x10 GbE SFP+ PCIe Adapter	1
90Y9430	3m Passive DAC SFP+ Cable	2
00FK932	System x 750W High Efficiency Platinum AC Power Supply	1
90Y3901	Integrated Management Module Advanced Upgrade	1
00FK622	System x Enterprise 2U Cable Management Arm (CMA)	1

Below is the bill of materials for a single Lenovo ThinkServer RD650 storage server.

Part #	Description	Quantity
70D00025UX	ThinkServer RD650: 2U Rack Server - 1 x Intel Xeon E5-2630 v3 2.40 GHz – 1X8 GB memory, 12Gb/s SAS 720ix AnyRAID controller, 1100W p/s	1
4X20F28577	ThinkServer Gen 5 1100W Platinum Hot Swap Power Supply	1
4XG0F28818	ThinkServer RD650 Intel Xeon E5-2630 v3 (8C, 85W, 2.4GHz) Processor	1
4X70F28589	ThinkServer 8GB DDR4-2133MHz (1Rx4) RDIMM	7
4XB0G88731	ThinkServer Gen 5 3.5" 4TB 7.2K Enterprise SAS 12Gbps HS HDD	12
4XC0F28742	Intel X520-DA2 AnyFabric 10Gb 2 Port SFP+ Ethernet Adapter	1
4XB0G45738	ThinkServer Gen 5 2.5" 300GB Value Read-Optimized SATA 6Gbps HS SSD	2
4XF0G45877	ThinkServer Gen 5 2.5" 2-Drive Rear Backplane Kit (Business Partner only)	1

`

3.2 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

4 Resources

For more information, see the following resources:

- NooBaa resources: NooBaa.com
- NooBaa Technical Review white paper: contact info@noobaa.com

Trademarks and special notices

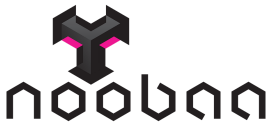

publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should 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.