Lenovo Big Data Reference Design for Cloudera Data Platform on ThinkSystem Servers

Last update: **June 2025**
Version 4.4

Reference architecture for Cloudera Data Platform with Apache Hadoop and Apache Spark

Solution based on the ThinkSystem servers

Deployment considerations for scalable racks including detailed validated bills of material

Faster and more accurate data analytics using GPU-enabled NVIDIA RAPIDS Accelerator for Apache Spark

Bill Zheng Sun

Ajay Dholakia

Table of Contents

1 Introduction

This document describes the reference design for Cloudera Data Platform (CDP) software on ThinkSystem servers. It provides architecture guidance for designing optimized hardware infrastructure for the Cloudera Data Platform Private Cloud edition, a distribution of Apache Hadoop and Apache Spark with enterprise-ready capabilities from Cloudera. This reference design provides the planning, design considerations, and best practices for implementing Cloudera Data Platform with Lenovo products.

Lenovo and Cloudera worked together on this document, and the reference architecture that is described herein was validated by Lenovo and Cloudera. The use of GPU-acceleration for Apache Spark 3 was validated by engineers at Lenovo, Cloudera and NVIDIA.

With the ever-increasing volume, variety and velocity of data becoming available to an enterprise comes the challenge of deriving the most value from it. This task requires the use of suitable data processing and management software running on a tuned hardware platform. With Apache Hadoop and Apache Spark emerging as popular big data storage and processing frameworks, enterprises are building so-called Data Lakes by employing these components.

Cloudera brings the power of Hadoop to the customer's enterprise. Hadoop is an open source software framework that is used to reliably manage large volumes of structured and unstructured data. Cloudera expands and enhances this technology to withstand the demands of your enterprise, adding management, security, governance, and analytics features. The result is that you get a more enterprise ready solution for complex, large-scale analytics.

The intended audience for this reference architecture is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing the big data solution with Lenovo hardware. It is assumed that you are familiar with Hadoop components and capabilities. For more information about Hadoop, see “Resources” on page 34.

2 Business problem and business value

Business Problem

The world is well on its way to generate 175 ZB of data by 2025. This is a 61% CAGR compare to 33 ZB of data in 2018. This data comes from everywhere, including sensors that are used to gather climate information, posts to social media sites, digital pictures and videos, purchase transaction records, and cell phone global positioning system (GPS) signals. This data is big data.

Big data spans the following dimensions:

- Volume: Big data comes in one size: large – in size, quantity and/or scale. Enterprises are awash with data, easily amassing terabytes and even petabytes of information.
- Velocity: Often time-sensitive, big data must be used as it is streaming into the enterprise to maximize its value to the business.
- Variety: Big data extends beyond structured data, including unstructured data of all varieties, such as text, audio, video, click streams, and log files.

Enterprises are incorporating large data lakes into their IT architecture to store all their data. The expectation is that ready access to all the available data can lead to higher quality of insights obtained through the use of analytics, which in turn drive better business decisions. A key challenge faced today by these enterprises is setting up an easy to deploy data storage and processing infrastructure that can start to deliver the promised value in a very short amount of time. Spending months of time and hiring dozens of skilled engineers to piece together a data management environment is very costly and often leads to frustration from unrealized goals. Furthermore, the data processing infrastructure needs to be easily scalable in addition to achieving desired performance and reliability objectives.

The need for and types of data analytics applications in an enterprise continues to expand. Collection, storage and processing of structured and unstructured data demands end-to-end data platform that can support a broad range of use cases such as archival storage, general processing using both standard and advanced in-memory computing algorithms, accelerated analytics like machine learning and deep learning, data warehousing with targeted support for analytics processing, real-time streaming and inference processing. This ever-growing list of target use cases demand a flexible and performant big data cluster design that can be optimized to address the set of business challenges at any given time.

Business Value

Lenovo has been delivering validated designs for big data clusters that combine CDP (and its predecessors) with targeted Lenovo systems to enable a whole host of big data use cases. Lenovo continues collaborating with industry-leading partners to create data processing and data management solutions suitable for various stages of the data flow. These solutions include:

- Data ingest and storage repositories for structured and unstructured data
- Compute clusters for machine learning and deep learning model training
- Tools for deploying trained models for inference across data center and edge IT infrastructure.

Big data is more than a challenge; it is an opportunity to find insight into new and emerging types of data to make your business more agile. Big data also is an opportunity to answer questions that, in the past, were beyond reach. CDP uses the latest big data technologies such as the in-memory processing capabilities of

Spark and Spark with GPU-enabled computing in addition to the standard MapReduce scale-out capabilities of Hadoop, to open the door to a world of possibilities.

How can businesses process tremendous amounts of raw data in an efficient and timely manner to gain actionable insights? CDP allows organizations to run large-scale, distributed analytics jobs on clusters of cost-effective server hardware. This infrastructure can be used to tackle large data sets by breaking up the data into "chunks" and coordinating data processing across a massively parallel environment. After the raw data is stored across the nodes of a distributed cluster, queries and analysis of the data can be handled efficiently, with dynamic interpretation of the data formatted at read time. The bottom line: Businesses can finally get their arms around massive amounts of untapped data and mine that data for valuable insights in a more efficient, optimized, and scalable way.

Hadoop is an open source software framework that is used to reliably manage and analyze large volumes of structured and unstructured data. Cloudera enhances this technology to withstand the demands of your enterprise, adding management, security, governance, and analytics features. The result is that you get an enterprise-ready solution for complex, large-scale analytics.

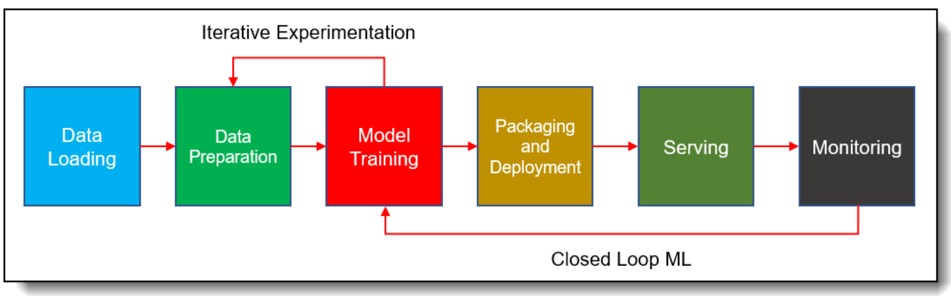

Figure 1. ***End-to-end data processing workflow***

Lenovo partners NVIDIA and Cloudera have integrated NVIDIA RAPIDS accelerated data science libraries on Cloudera Data Platform (CDP), enabling GPU-accelerated Apache Spark 3 applications. Apache Spark included in CDP has been a workhorse for numerous data analytics tasks such as batch/real-time streaming, data warehouse and machine learning among others. Accelerating Spark with GPU-enabled computation is the next leap forward in helping enterprises achieve the goal of faster and better model development. Key technical features include faster and more accurate analytics and prediction from machine learning models. The ultimate business benefits are significantly improved price-to-performance as well as return-on-investment (ROI) by leveraging better data analytics.

CDP deployed on Lenovo ThinkSystem servers provides excellent price/performance. The reference architecture supports entry through high-end configurations and the ability to easily scale as the use of big data grows. A choice of infrastructure components provides flexibility in meeting varying big data analytics requirements.

3 Requirements

The functional and non-functional requirements for this reference architecture are desribed in this section.

3.1 Functional Requirements

A big data solution supports the following key functional requirements:

- Ability to handle various workloads, including batch and real-time analytics
- Industry-standard interfaces so that applications can work with Cloudera
- Ability to handle large volumes of data of various data types
- Various client interfaces

3.2 Non-functional Requirements

Customers require their big data solution to be easy, dependable, and fast. The following non-functional requirements are key:

- Easy:
    - Ease of development
    - Easy management at scale
    - Advanced job management
    - Multi-tenancy
    - Easy to access data by various user types
- Dependable:
    - Data protection with snapshot and mirroring
    - Automated self-healing
    - Insight into software/hardware health and issues
    - High availability (HA) and business continuity
- Fast:
    - Superior performance
    - Scalability
- Secure and governed:
    - Strong authentication and authorization
    - Kerberos support
    - Data confidentiality and integrity

4 Architectural Overview

4.1 Cloudera Data Platform

Figure 2 shows the high-level architecture of the Cloudera data platform that consists of CDP private cloud built on Lenovo hardware and CDP public cloud running on public clouds. CDP private cloud is built for hybrid cloud, seamlessly connecting on-premises environments to public clouds with consistent, built-in security and governance. CDP Public Cloud is a cloud form factor of CDP. Cloudera SDX (Shared Data Experience) secure and govern platform data and metadata as well as control capabilities. Data security, governance, and control policies are set once and consistently enforced everywhere, reducing operational costs and business risks while also enabling complete infrastructure choice and flexibility.

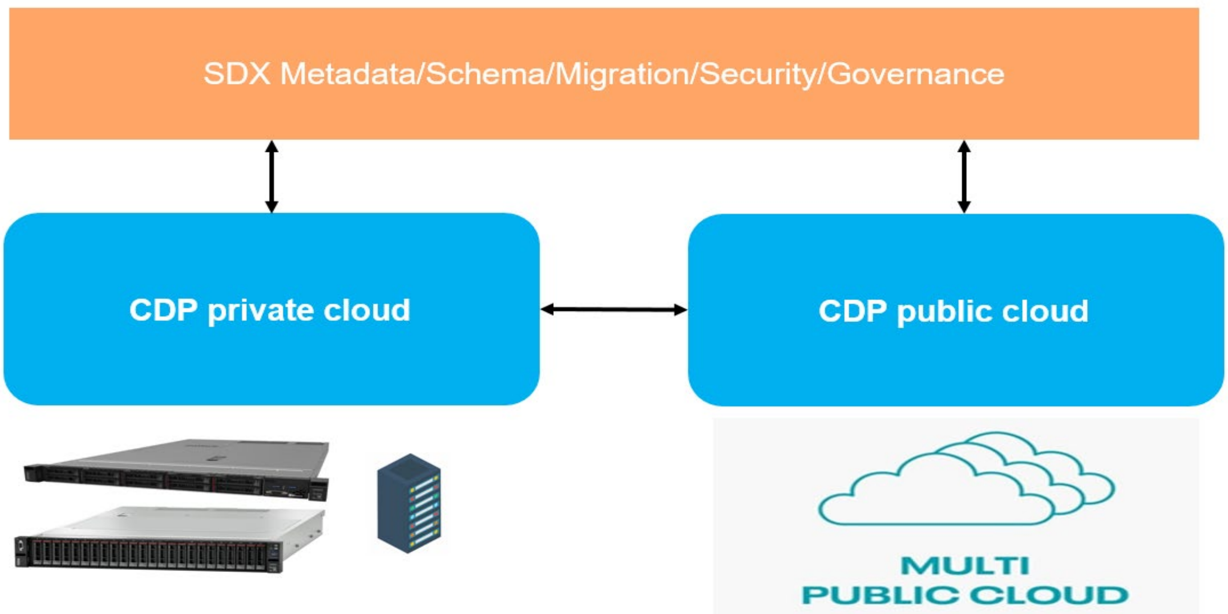

Figure 2. Cloudera Data Platform architecture overview

4.2 CDP Private Cloud

Cloudera Data Platform (CDP) Private Cloud is the newest on-premises version of CDP that brings many benefits of the public cloud services to the on-premises deployment.

CDP Private Cloud provides a disaggregation of compute and storage, and allows independent scaling of compute and storage clusters. CDP Private Cloud gets unified security, governance, and metadata management through Cloudera SDX.

CDP Private Cloud users can rapidly provision and deploy Cloudera Data Warehousing and Cloudera Machine Learning services through the Management Console, and easily scale them up or down as required.

Figure 3 shows a CDP Private Cloud deployment. It requires you to have a Private Cloud Base cluster and a Private Cloud Experiences cluster deployed on a RedHat OpenShift cluster. Both Private Cloud Base cluster and the OpenShift cluster are set up on Lenovo servers. The Private Cloud deployment process involves configuring Management Console on the OpenShift cluster, registering an environment by providing details of the Data Lake configured on the Base cluster, and then creating the workloads.

The Cloudera Private Cloud Base is comprised of a variety of components such as Apache HDFS, Apache OZone, Apache Spark, and Apache Impala, along with many other components for specialized workloads. You can select any combination of these services to create clusters that address your business requirements and workloads. CDP Private Cloud Base solutions described in this document can be deployed on bare-metal infrastructure. This means that both the management nodes and the data nodes are implemented on physical hosts. The number of servers of each type is determined based on requirements for high-availability, total data capacity and desired performance objectives. This reference design provides validated solutions for traditional local storage on the Lenovo server, configured for non-RAID JBOD (Just-a-Bunch-Of-Drives) which gives over 40% more storage capacity per rack and more compute nodes compared to nodes with internal HDD storage.

With Hadoop local storage, Lenovo server contains compute and storage in the same physical enclosure. Scale out is accomplished by adding one or more nodes which add both compute and storage simultaneously to the cluster. The Lenovo server provides the highest CPU core count and highest total memory per node for a very high-end analytics solution.

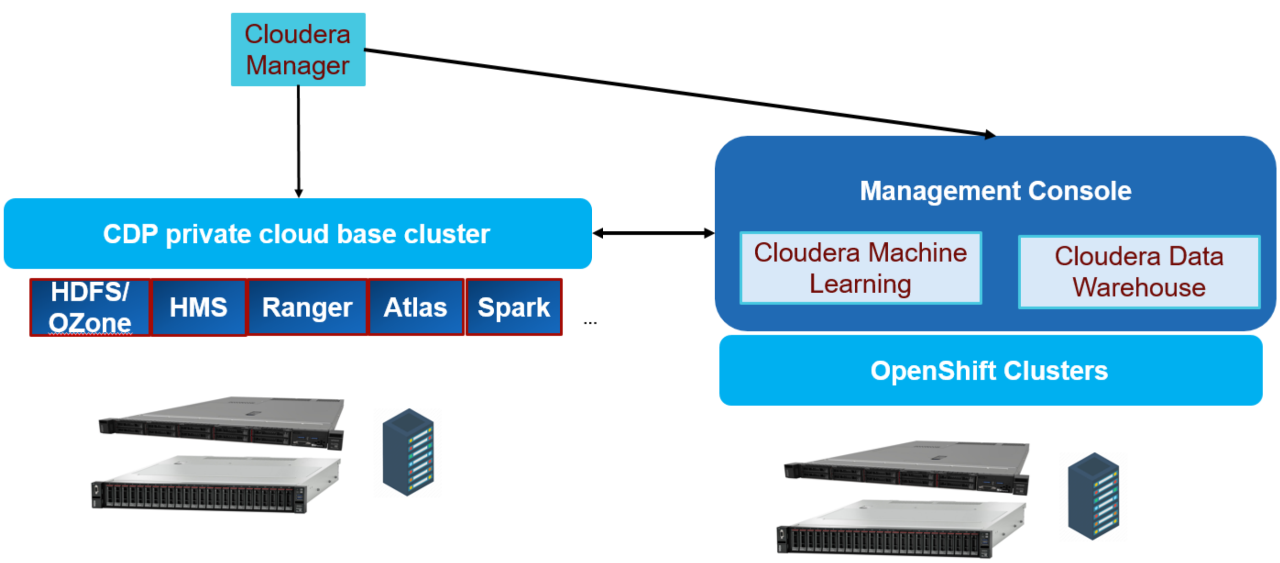

Figure 3. CDP Private Cloud

Cloudera Private Cloud base provides several interfaces that allow administrators and users to perform administration and data functions, depending on their roles and access level. Hadoop application programming interfaces (APIs) can be used to access data. Cloudera APIs can be used for cluster management and monitoring. Cloudera data services, management services, and other services run on the nodes in cluster. Storage is a component of each data node in the cluster. Data can be incorporated into Cloudera Data Platform storage through the Hadoop APIs, depending on the needs of the customer.

5 Component Model

Cloudera Data Platform provides features and capabilities that meet the functional and nonfunctional requirements of customers. It supports mission-critical and real-time big data analytics across different industries, such as financial services, retail, media, healthcare, manufacturing, telecommunications, government organizations, and leading Fortune 100 and Web 2.0 companies.

Cloudera Data Platform is the world's most complete, tested, and popular distribution of Apache Hadoop and related projects. All of the packaging and integration work is done for you, and the entire solution is thoroughly tested and fully documented. By taking the guesswork out of building out your Hadoop deployment, Cloudera Data Platform gives you a streamlined path to success in solving real business problems with big data.

The Cloudera platform for big data can be used for various use cases from batch applications that use MapReduce or Spark with data sources, such as click streams, to real-time applications that use sensor data.

Figure 4 shows the Cloudera Data Platform that meet the functional requirements of customers.

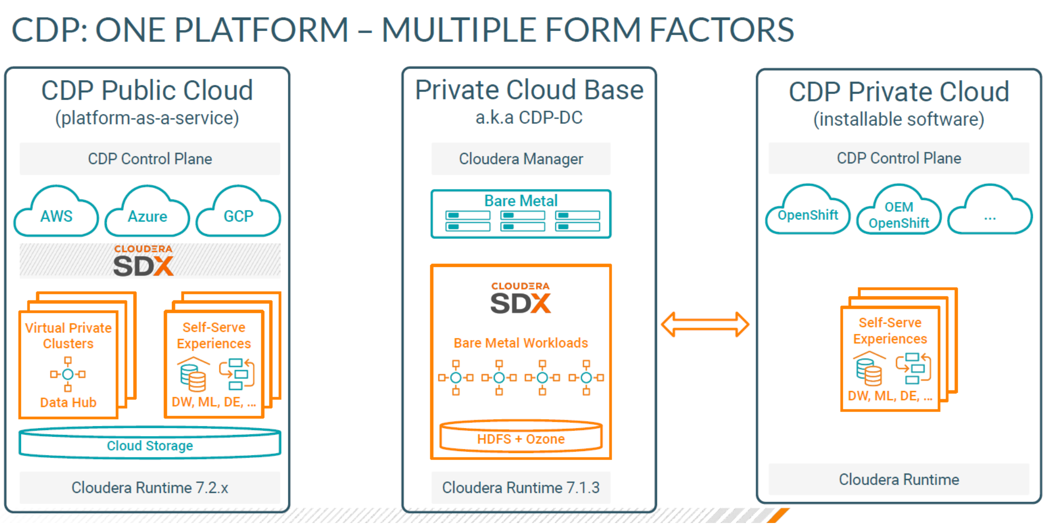

Figure 4. Cloudera Data Platform overview

5.1 Cloudera Components

Cloudera Data Platform solution contains the following components:

- Analytic SQL: Apache Impala

  Impala is the industry's leading massively parallel processing (MPP) SQL query engine that runs natively in Hadoop. Apache-licensed, open source Impala project combines modern, scalable parallel database technology with the power of Hadoop, enabling users to directly query data stored in HDFS and Apache HBase without requiring data movement or transformation. Impala is designed from the ground up as part of the Hadoop system and shares the same flexible file and data formats, metadata, security, and resource management frameworks that are used by MapReduce, Apache Hive and other components of the Hadoop stack.

- Search Engine: Cloudera Search

  Cloudera Search is Apache Solr fully integrated in the Cloudera platform, taking advantage of the flexible, scalable, and robust storage system and data processing frameworks included in Cloudera Data Platform (CDP). This eliminates the need to move large data sets across infrastructures to perform business tasks. It further enables a streamlined data pipeline, where search and text matching is part of a larger workflow. Cloudera Search also includes valuable integrations that make searching more scalable, easy to use, and optimized for near-real-time and batch-oriented indexing. These integrations include Cloudera Morphlines, which is a customizable transformation chain that simplifies loading any type of data into Cloudera Search.

- NoSQL - HBase

  A scalable, distributed column-oriented datastore. HBase provides real-time read/write random access to very large datasets hosted on HDFS.

- Stream Processing: Apache Spark

  Apache Spark is an open source, parallel data processing framework that complements Hadoop to make it easy to develop fast, unified big data applications that combine batch, streaming, and interactive analytics on all your data. Cloudera offers commercial support for Spark with Cloudera Data Platform. Spark is 10 – 100 times faster than MapReduce which delivers faster time to insight, allows inclusion of more data, and results in better business decisions and user outcomes.

- Machine Learning: Spark MLlib

  MLlib is the API that implements common machine learning algorithms. MLlib is usable in Java, Scala, Python and R. Leveraging Spark's excellence in iterative computation, MLlib runs very fast, high-quality algorithms.

- Cloudera Manager

  Cloudera Manager is the industry's first and most sophisticated management application for Hadoop and the enterprise data hub. Cloudera Manager sets the standard for enterprise deployment by delivering granular visibility into and control over every part of the data hub, which empowers operators to improve performance, enhance quality of service, increase compliance, and reduce administrative costs. Cloudera Manager makes administration of your enterprise data hub simple and straightforward, at any scale. With Cloudera Manager, you can easily deploy and centrally operate the complete big data stack.

Cloudera Manager automates the installation process, which reduces deployment time from weeks to minutes; gives you a cluster-wide, real-time view of nodes and services running; provides a single, central console to enact configuration changes across your cluster; and incorporates a full range of reporting and diagnostic tools to help you optimize performance and utilization.

- Cloudera Manager Metrics

  Cloudera Manager monitors several performance metrics for services and role instances that are running on your clusters. These metrics are monitored against configurable thresholds and can be used to indicate whether a host is functioning as expected. You can view these metrics in the Cloudera Manager Admin Console, which displays metrics about your jobs (such as the number of currently running jobs and their CPU or memory usage), Hadoop services (such as the average HDFS I/O latency and number of concurrent jobs), your clusters (such as average CPU load across all your hosts) and so on.

- Replication Manager
  Cloudera Manager provides an integrated, easy-to-use management solution to enable data protection on the Hadoop platform. Replication Manager enables you to replicate data across data centers for disaster recovery scenarios.
  Replications can include data stored in HDFS, data stored in Hive tables, Hive metastore data, and Impala metadata (catalog server metadata) associated with Impala tables registered in the Hive metastore. When critical data is stored on HDFS, Cloudera Manager helps to ensure that the data is always available, even in case of complete shutdown of a data center. You can also use the HBase shell to replicate HBase data. (Cloudera Manager does not manage HBase replications.)
- Cloudera Manager API

  The Cloudera Manager API provides configuration and service lifecycle management, service health information and metrics, and allows you to configure Cloudera Manager. The API is served on the same host and port as the Cloudera Manager Admin Console, and does not require an extra process or extra configuration. The API supports HTTP Basic Authentication, accepting the same users and credentials as the Cloudera Manager Admin Console.

- Cloudera Kafka

  Cloudera Distribution of Apache Kafka is a distributed commit log service. Kafka functions much like a publish/subscribe messaging system, but with better throughput, built-in partitioning, replication, and fault tolerance. Kafka is a good solution for large scale message processing applications. It is often used in tandem with Apache Hadoop and Spark Streaming.

- Data Warehouse
  The CDW Private Cloud service enables self-serve creation of independent data warehouses and data marts for teams of business analysts without the overhead of bare metal deployments. In the CDW Private Cloud service, your data is stored in HDFS in the base cluster. The service is composed of:

  - Database Catalogs:

A logical collection of metadata definitions for managed data with its associated data context. The data context is comprised of table and view definitions, transient user and workload contexts from the Virtual Warehouse, security permissions, and governance artifacts that support functions such as auditing. One Database Catalog can be queried by multiple Virtual Warehouses.

  - Virtual Warehouses:

    An instance of compute resources that is equivalent to a cluster. A Virtual Warehouse provides access to the data in tables and views that correlate to a specific Database Catalog. Virtual Warehouses bind compute and storage by executing queries on tables and views that are accessible through the Database Catalog that they have been configured to access.

- Cloudera Machine Learning
  Machine learning has become one of the most critical capabilities for modern businesses to grow and stay competitive today. From automating internal processes to optimizing the design, creation, and marketing processes behind virtually every product consumed, ML models have permeated almost every aspect of our work and personal lives.

  Cloudera Machine Learning (CML) is Cloudera's new cloud-native machine learning service, built for CDP. The CML service provisions clusters, also known as ML workspaces, that run natively on Kubernetes.

  Each ML workspace enable teams of data scientists to develop, test, train, and ultimately deploy machine learning models for building predictive applications all on the data under management within the enterprise data cloud. ML workspaces are ephemeral, allowing you to create and delete them on-demand. ML workspaces support fully containerized execution of Python, R, Scala, and Spark workloads through flexible and extensible engines.

- Management Console
  The Management Console is a service used by CDP administrators to manage environments, users, and services.

  The Management Console allows you to:

  - Enable user access to CDP Private Cloud experiences, onboard and set up authentication for users, and determine access rights for the various users to the available resources.
  - Register an environment, which represents the association of your user account with compute resources using which you can manage and provision workloads such as Data Warehouse and Machine Learning. When registering the environment, you must specify a Data Lake residing on the Private Cloud base cluster to provide security and governance for the workloads.
  - View information about the resources consumed by the workloads for an environment.
  - Collect diagnostic information from the services for troubleshooting purposes.

For more information, see this website: cloudera.com/content/cloudera/en/products-and-services/product-comparison.html

The Cloudera solution is operating system independent. Cloudera supports many Linux® operating systems, including Red Hat Linux and CentOS. For more information about the versions of supported operating systems, see this website:

https://docs.cloudera.com/cdp-private-cloud/latest/release-guide/topics/cdpdc-os-requirements.html

5.2 Apache Spark 3

Apache Spark is an open source, parallel data processing framework for big data workloads. It uses the in-memory cache and optimized query execution mode to quickly analyze and query data of any size. It provides development APIs using Java, Scala, Python, and R languages, and supports code reuse across multiple workloads-batch processing, interactive queries, real-time analysis, machine learning, and graphics processing. In Cloudera CDP, it complements Hadoop/Ozone to make it easy to develop fast, unified big data applications on all your data.

Spark applications can run in a distributed mode on a cluster using a master/slave architecture that uses a central coordinator called “driver” and potentially large number of “worker” processes that execute individual tasks in a Spark job. The Spark executor processes also provide reliable in-memory storage of data distributed across the various nodes in a cluster. The components of a distributed Spark application are shown in Figure 5

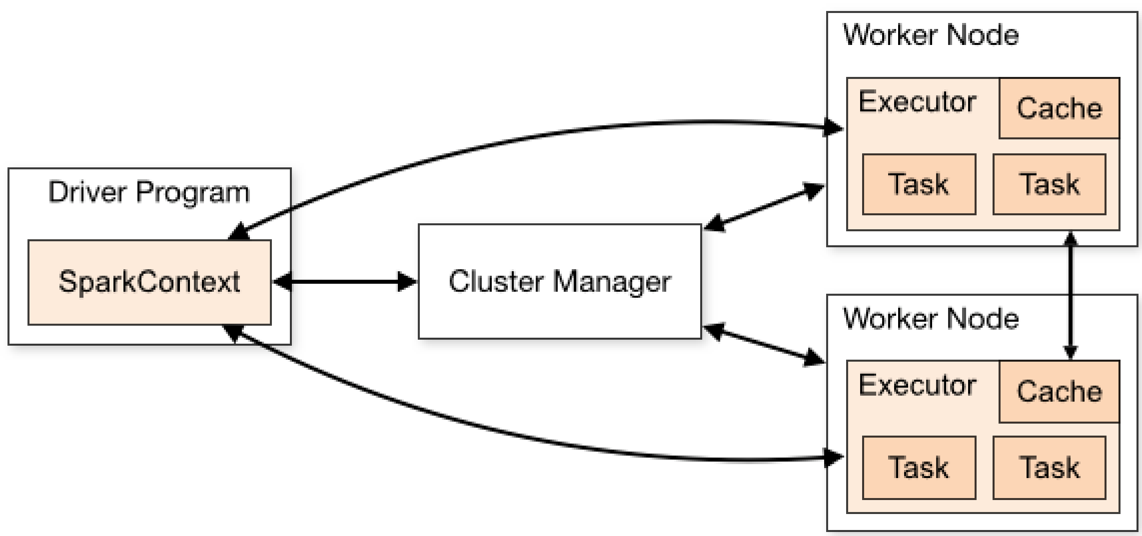

Figure 5. Distributed Spark application component model

A key distinguishing feature of Spark is the data model, based on RDDs (Resilient Distributed Datasets). This model enables a compact and reusable organization of data set that can reside in the main memory and can be accessed by multiple tasks. Iterative processing algorithms can benefit from this feature by not having to store and retrieve data sets from disks between iterations of computation. These capabilities are what deliver the significant performance gains compared to MapReduce.

RDDs support two types of operations: Transformations and Actions. Transformations are operations that return a new RDD, while Actions return a result to the driver program. Spark groups operations together to reduce the number of passes taken over the data. This so-called lazy evaluation technique enables faster data processing. Spark also allows caching data in memory for persistence to enable multiple uses of the same data. This is another technique contributing to faster data processing.

Cloudera offers commercial support for Spark with Cloudera Data Platform. Spark is 10-100 times faster than MapReduce which delivers faster time to insight, allows inclusion of more data, and results in better business decisions and user outcomes. First version of Apache Spark 3.x(Spark 3.0) was officially released on June 18, 2020. Apache Spark 3.0 has brought significant progress in Python and SQL functions and focused on ease of development and production. The performance of first version of Apache Spark 3.x (Spark3.0) in the entire runtime is about twice fast than Spark 2.4.

Here are main new features of Spark3.0:

1. In the TPC-DS benchmark test, by enabling other optimization measures such as adaptive query execution, dynamic partition pruning, etc., compared to Spark 2.4, the performance is improved by 2 times

2. Compatible with ANSI SQL

3. Major improvements to the pandas API, including python type hints and other pandas UDFs

4. Simplify Pyspark exceptions and better handle Python errors

5. New UI for structured streaming

6. The speed of calling R language UDF is increased by 40 times

Spark mitigated the I/O problems found in MapReduce, but the bottleneck has shifted from I/O to compute for a growing number of applications. Spark can have a several times performance improvement with the advent of GPU-accelerated computing.

NVIDIA RAPIDS is a suite of open-source software libraries and API for executing end-to-end data science and analytics pipelines entirely on GPUs. RAPIDS is built on top of NVIDIA CUDA, an architecture and software platform for GPU computing. RAPIDS expose GPU parallelism and high-bandwidth memory speed through user-friendly APIs. The way RAPIDS accelerates computing is with integration to leading distributed computing frameworks like Apache Spark and data science frameworks, such as Pytorch and TensorFlow.

The RAPIDS Accelerator for Apache Spark takes a different approach to accelerating data science workloads on GPUs. Fundamentally, its approach is to provide transparent acceleration of Spark data frame jobs via a Spark plugin that integrates with Spark's query planner. The plugin rewrites data frame query plans in order to evaluate accelerable operations with implementations that use libcudf (the C++ library providing accelerated data frame functionality to the Python cuDF library) to execute on the GPU. Operations that cannot be accelerated will run on the CPU with Spark's built-in implementations; if there is a branch of a query plan that includes both accelerable and non-accelerable operations, the RAPIDS Accelerator plugin will automatically insert transfers between host and device memory so that both kinds of operations can work together transparently to execute a given query plan.

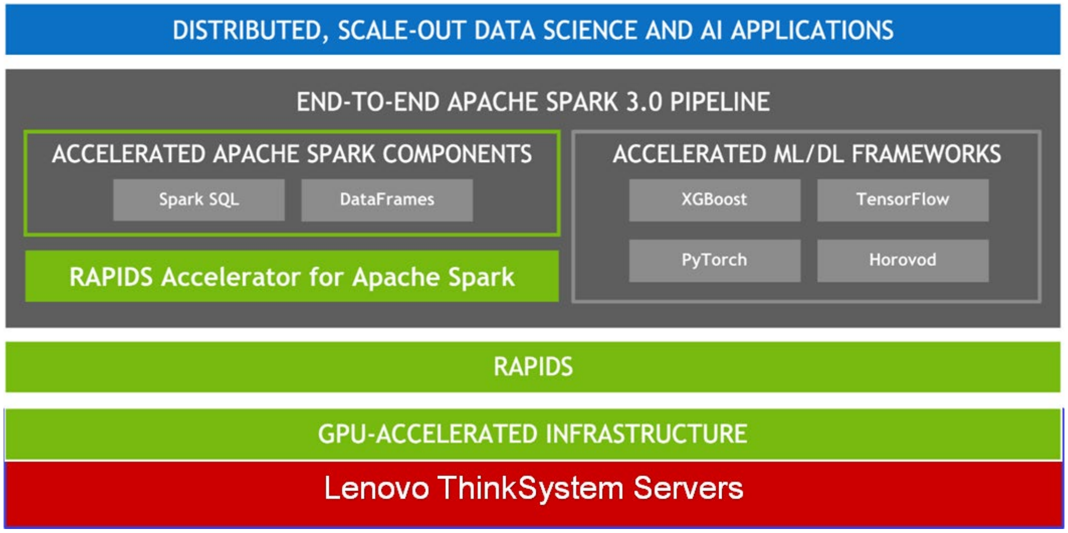

Figure 6. Apache Spark 3.x GPU-Accelerated stack on Lenovo ThinkSystem Servers

The RAPIDS Accelerator for Apache Spark also provides an accelerated shuffle implementation using UCX (for data transfer within clusters) and integration with GPU-accelerated XGBoost.

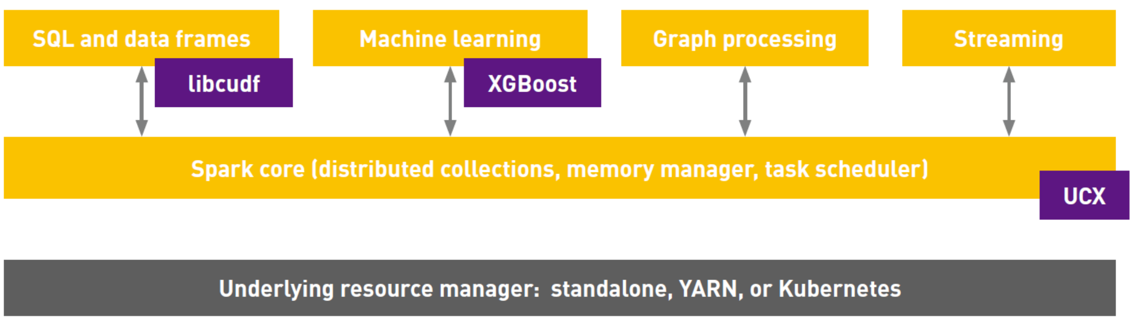

Figure 7. The architecture of GPU-accelerated Spark

RAPIDS takes advantage of Apache Arrow. It offers a powerful GPU DataFrame that is compatible with Apache Arrow data structures with a familiar DataFrame API. Arrow defines a standardized language-independent columnar memory format optimized for data locality to accelerate analytical processing performance on modern hardware like CPUs and GPUs. It provides zero-copy data reads for streaming messaging and inter-process communication without serialization overhead.

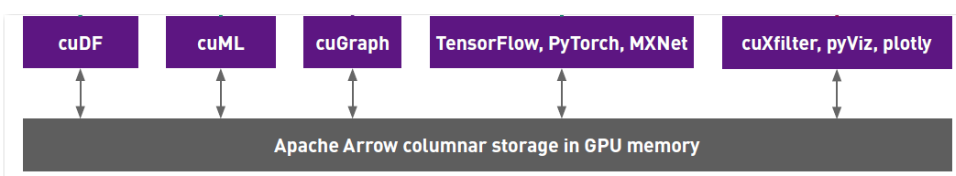

Figure 8. Data processing with Apache Arrow

6 Operational Model

This section describes the operational model for the Cloudera reference architecture. To show the operational model for different sized customer environments, HW selection and considerations are discussed in this section.

A Cloudera deployment consists of cluster nodes, networking equipment, power distribution units, and racks. The predefined configurations can be implemented as-is or modified based on specific customer requirements, such as lower cost, improved performance, and increased reliability. Key workload requirements, such as the data growth rate, sizes of datasets, and data ingest patterns help in determining the proper configuration for a specific deployment. A best practice when a Cloudera cluster infrastructure is designed is to conduct the proof of concept testing by using representative data and workloads to ensure that the proposed design works.

6.1 Lenovo Server Description

Lenovo ThinkSystem servers are ideal rack servers 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. Lenovo ThinkSystem servers are designed to handle a wide range of workloads, such as big data, databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), infrastructure security, systems management, enterprise applications, collaboration/email, streaming media, web, and HPC. In particular, big data workloads require servers with scalable performance and storage capacity and a feature set that can enable flexible configurations to meet the needs of a variety of node types in a cluster.

6.1.1 Server Selection

Lenovo ThinkSystem server portfolio offers multiple rack unit sizes and processor socket counts featuring the latest generation of Intel and AMD processors. The choice of servers for big data clusters is driven by a variety of factors including target use cases, performance, flexible storage capacity, I/O and acceleration support as well as cost. It is recommended to engage with your Lenovo contacts to help determined the desired servers for your collection of intended use cases.

- **Lenovo ThinkSystem SR650 V3 and SR630 V3**: The SR650 V3 and SR630 V3 are based on the Intel 5th Gen Xeon Scalable processor family (formerly codenamed " Emeral Rapids "), or Intel 4th Gen Xeon Scalable processor family (formerly codenamed " Sapphire Rapids "). These servers have been designed to take advantage of the features of the 5th generation Intel Xeon Scalable processors, able to scale up to 64 cores on the highest end 385W SKU with GPU support for 8x single or 3x double width.

  Intel's new 5th generation Intel Xeon Scalable processors (formerly codenamed " Emerald Rapids ") introduce significant performance enhancements. Emerald Rapids processors support up to 64 cores per CPU while 60 cores in Sapphire Rapids, marking a nearly 3x the maximum last-level cache compared to the previous generation Sapphire Rapids CPU. They offer eight channels of DDR5 per CPU and improved inter-socket bandwidth. There is about 30%+ boost in average performance per watt across a range of workloads, making the Emerald Rapids processors a notable advancement in computational efficiency.

The SR650 V3 is a configuration-rich offering, supporting 30 different drive bay configurations in the front, middle and rear of the server and 5 different slot configurations at the rear of the server. This server is two rack units in height and is ideally suited for providing various worker node configurations. The SR630 V3 is one rack unit in height and is ideally suited for master and utility node configurations. This server also offers onboard NVMe PCIe ports that allow direct connections to 16x NVMe SSDs, which results in faster access to store and access data.

- **Lenovo ThinkSystem SR650 V4:** The SR650 V4 server is 2-socket 2U rack server designed for enterprise workloads requiring scalability, flexibility, and energy efficiency. It is based on Intel Xeon 6700-series or Xeon 6500-series processors, with Performance-cores (P-cores), formerly codenamed "Granite Rapids-SP". It supports up to 86 cores and 172 threads with core speeds of up to 4GHz and TDP ratings of up to 350W. It can support up to 32 DDR5 memory DIMMs, supporting speeds up to 6400 MHz. Using 256GB 3D RDIMMs, the server supports up to 8TB of system memory. By supporting CXL 2.0 standard, the server is well-suited for CPU-memory intensive workload in large-scale data processing. The server supports up to 10 single-width GPUs or 2 double-width GPUs, catering to high-performance computing needs.

- **Lenovo ThinkSystem SR630 V4**: The SR630 V4 server is 2-socket 1U rack server based on Intel Xeon 6700-series or Xeon 6500-series processors, with Performance-cores (P-cores), formerly codenamed "Granite Rapids-SP" or with Efficient-cores (E-cores), formerly codenamed "Sierra Forest-SP". The server shipped with Performance-cores (P-cores) can support up to 86 cores and 172 threads with core speeds of up to 4GHz and TDP ratings of up to 350W. The server shipped with Efficient-cores (E-cores) supports Up to 144 cores with core speeds of up to 2.4 GHz and TDP ratings of up to 330 W. Supporting CXL 2.0 make this server well-suited for next-generation CPU-memory intensive workloads, with reduced compute latency lower TCO. It supports up to three single-width GPUs each up to 75W for processing power in a 1U system.

  Both ThinkSystem SR650 V4 and SR630 V4 servers are shipped with XClarity Controller 3 (XCC3) to monitor server availability. Optional upgrade to XCC3 Premier can offer additional features such as Neighbour Groups, System Guard, a CNSA-compliant security mode, a FIPS 140-3-compliant mode, and enhanced NIST 800-193 support.

- **Lenovo ThinkSystem SR665 and SR645**: The SR665 and SR645 servers have been designed to take advantage of the features of the AMD EPYC 7002 and AMD EPYC 7003 processors, such as the full performance of 280W 64-core processors, support for 3200 MHz memory and PCIe Gen 4.0 support. These servers are 2-sock servers that feature the AMD EPYC 7002 "Rome" and AMD EPYC 7003 "Milan" families of processors. With up to 64 cores per processor and support for the new PCIe 4.0 standard for I/O, these servers are well-suited for dense workloads that can take advantage of GPU processing and high-performance NVMe drives. The SR665 offers the ultimate in two-socket server performance in a 2U form factor. In particular, it is ideally suited for various worker node configurations in a big data cluster. The SR645 offers the ultimate in two-socket server performance in a space-saving 1U form factor and is ideally suited for master and utility node configurations.

- **Lenovo ThinkSystem SR665 V3 and SR645 V3**: The Lenovo ThinkSystem SR665 V3 and SR645 V3 are 2-socket servers that features the 5th GEN AMD EPYC 9005 "Turin" family of processors. With up to 160 “Zen 5” cores per processor with each core speed of up to 4.0 GHz, these servers support up to 24 DDR5 memory DIMMS, 12 memory channels per processor (1DIMM per channel),128 PCIe 5.0 I/O lanes with 64 lanes available for PCIe and NVMe devices, well-suited for dense workloads that can take advantage of GPU processing and high-performance NVMe drives. The new “Zen 5” core architecture, provides up to 17% better instructions per clock (IPC) for enterprise and cloud workloads and up to 37% higher IPC in AI and high performance computing (HPC) compared to “Zen 4”. ThinkSystem SR665 V3 offers the ultimate in two-socket server performance in a 2U form factor. ideally suited for various worker node configurations in a big data cluster. ThinkSystem SR645 V3 is a dense, high performance, 2-socket 1U rack server, suitable for running as master and utility node for CDP deployment.

6.1.2 Processor Selection

Minimum Hadoop recommendations are 1 CPU processor core per data disk plus additional cores dedicated to specific Cloudera software services and data analytics functions. The worker node processors in this reference architecture provide 1 processor cores per HDD ratio which gives the maximum HDD throughput plus a full set of cores for additional data analytics.

Cloudera workload types may be skewed toward IO-bound workloads that create heavy network traffic or CPU bound workloads that stress the CPU cores themselves. Processors provide higher core counts to meet the highest of CPU bound workloads.

Below are several examples of IO-bound workloads:

- Sorting

- Indexing
- Grouping
- Data importing and exporting
- Data movement and transformation

Below are several examples of CPU-bound workloads:

- Clustering/Classification
- Complex text mining
- Natural-language processing
- Feature extraction

Moving up to the processor categories add higher total core counts, higher operating frequency (in Megahertz, MHz), and increased internal memory cache sizes. Processor cost increases incrementally as each of the processor specifications increase. Cloudera provides processor selection guidance for sizing based on specific types of Cloudera cluster services required. Reference this link below for details:

https://docs.cloudera.com/cdp-private-cloud-base/latest/installation/topics/cdpdc-hardware-requirements.html

6.1.3 GPU Accelerator Selection

Lenovo ThinkSystem servers are designed to include one or more GPUs that provide compute acceleration for demanding workloads.

NVIDIA H100 is powered by NVIDIA Hopper Architecture. It delivers acceleration at every scale for AI, data analytics, and HPC to tackle the world's toughest computing challenges. As the engine of the NVIDIA data center platform, H100 can efficiently scale up to thousands of GPUs or, using second-generation Multi-Instance GPU (MIG) technology, can be partitioned into up to seven isolated GPU instances to accelerate workloads of all sizes. H100's fourth-generation Tensor Core technology now accelerates more levels of precision for diverse workloads, speeding time to insight as well as time to market.

NVIDIA GPU with RAPIDS accelerator for Apache Spark can speed up bigdata processing by several times.

Other types of NVIDIA GPUs are also supported by Lenovo ThinkSystem Server. Following table shows GPUs families and their target workloads.

Table 1. GPU families and workloads

Form factor	NVIDIA Data Science and Virtualization	NVIDIA 3D Graphics
SXM	H100 SXM5 A100 SXM	

Dual slot	H100 & H100 NVL L40S L40 A100 A30 A16	A40 RTX A6000 RTX A4500 RTX A2000
Single slot	L4 A10 A2 Tesla T4	Quadro RTX T1000 Quadro RTX T400

GPUs in Data Science and virtualization column are suitable for Bigdata processing acceleration in Spark workload.

6.1.4 Memory Size and Performance

Low node memory capacity can negatively impact cluster performance by causing workload thrashing and spilling to slower storage devices. Also, in-memory workloads such as Apache Spark benefit from larger memory capacity. Spark workloads are recommended to use higher memory capacities than with Hadoop map reduce for this reason.

This reference architecture specifies node memory sizes appropriate to Lenovo server types and with moderate Cloudera workloads, but many memory choices are available. Table 2 shows memory capacity recommendations based on cluster node type.

Table 2. Node memory capacity recommendations

Memory Capacity	Node Type
192 GB	Master and Utility nodes
288 - 384 GB	Worker node minimum
578 - 3,000 GB	Worker node in-memory Spark and high-performance workloads

6.1.5 Estimating Disk Space

When you are estimating disk space within a Cloudera Data Platform cluster, consider the following:

- For improved fault tolerance and performance, Cloudera Data Platform replicates data blocks across multiple cluster worker nodes. By default, the file system maintains three replicas.
- Compression ratio is an important consideration in estimating disk space and can vary greatly based on file contents. If the customer’s data compression ratio is unavailable, assume a compression ratio of 2.5:1.

- To ensure efficient file system operation and to allow time to add more storage capacity to the cluster if necessary, reserve 25% of the total capacity of the cluster.

Assuming the default three replicas maintained by Cloudera Data Platform, the raw data disk space and the required number of nodes can be estimated by using the following equations:

Total raw data disk space = (User data, uncompressed) * (4 / compression ratio)

Total required worker nodes = (Total raw data disk space) / (Raw data disk per node)

You should also consider future growth requirements when estimating disk space.

Based on these sizing principals, Table 3 shows an example for a cluster that must store 500 TB of uncompressed user data. The example shows that the Cloudera cluster needs 800 TB of raw disk space to support 500 TB of uncompressed data. The 800 TB is for data storage and does not include operating system disk space. A total of 15 nodes are required to support a deployment of this size.

Total raw data disk space = 500TB * (4 / 2.5) = 500 * 1.6 = 800TB

Total required worker nodes = 800TB / (4TB * 20 drives) = 800TB / 80TB = 10 => 10 nodes

Table 3. Example of storage sizing with 4TB drives

Description	Value
Data storage size required (uncompressed)	500 TB
Compression ratio	2.5:1
Size of compressed data	200 TB
Storage multiplication factor	4
Raw data disk space needed for Cloudera cluster	800 TB
Storage needed for Cloudera Hadoop 3x replication	600 TB
Reserved storage for headroom (25% of 800TB)	200 TB
Raw data disk per node (with 4TB drives * 20 drives)	56 TB
Minimum number of nodes required (800/80)	10

6.2 Cluster Node Configurations

The Cloudera reference design architecture is implemented on a set of nodes that make up a cluster which includes two main node types: Worker nodes and Master nodes. Worker nodes use ThinkSystem 2U rack servers with locally attached storage. Master nodes use ThinkSystem 1U rack servers.

Worker nodes run data (worker) services for storing and processing data.

Master nodes run the following types of services:

- Management control services for coordinating and managing the cluster
- Miscellaneous and optional services for file and web serving

6.2.1 Worker Nodes

A balance in performance vs. cost for Cloudera worker nodes.   Higher core count and frequency processors can bring higher performance for compute intensive workloads. A minimum of 384 GB of memory is recommended for most MapReduce workloads with 768 GB or more recommended for HBase, Spark and memory-intensive MapReduce workloads.

The OS is loaded on a dual M.2 SSD memory module with RAID1 mirroring capability for High Availability and lowest cost.   Data disks are JBOD configured for maximum Hadoop and Spark performance with data fault tolerance coming from the HDFS file system 3x replication factor.

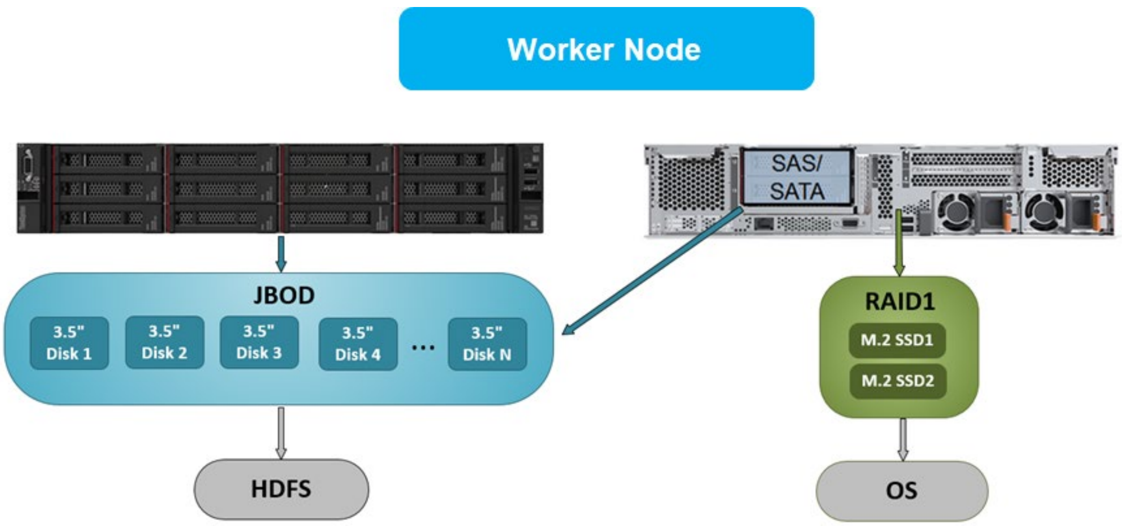

Figure 9. Worker node disk assignment

6.2.2 Master and Utility Nodes

The Master node is the nucleus of the Hadoop Distributed File System (HDFS) and supports several other key functions that are needed on a Cloudera cluster. Master nodes primarily run the HDFS NameNode and JournalNode service.   A utility node runs remaining management services including an additional instance of ZooKeeper. The Master node runs the following services:

YARN ResourceManager: Manages and arbitrates resources among all the applications in the system.

Hadoop NameNode: Controls the HDFS file system. The NameNode maintains the HDFS metadata, manages the directory tree of all files in the file system and tracks the location of the file data within the cluster. The NameNode does not store the data of these files.

ZooKeeper: Provides a distributed configuration service, a synchronization service and a name registry for distributed systems.

JournalNode: Collects, maintains and synchronize updates from NameNode.

HA ResourceManager: Standby ResourceManager that can be used to provide automated failover.

HA NameNode: Standby NameNode that can be used to provide automated failover.

Utility nodes run other non-master services for Hadoop management such as: Cloudera Manager, HBase master, HiveServer2, and Spark History Server.

The M.2 SSD form factor is intended for Operating System storage in this reference design architecture.

The Master node uses 10 drives for the following storage pools:

- Two drives (M.2 SSD modules) are configured with RAID 1 for operating system
- Two drives are configured with RAID 1 for NameNode metastore
- Four drives are configured with RAID 10 for database
- One drive is configured with RAID 0 for ZooKeeper
- One drive is configured with RAID 0 for Quorum Journal Node store

This design separates the data stores for different services and provides best performance. SSD drives in the 2.5" and 3.5" SAS/SATA form factor and PCIe card flash storage can be used to provide improved I/O performance for the database.

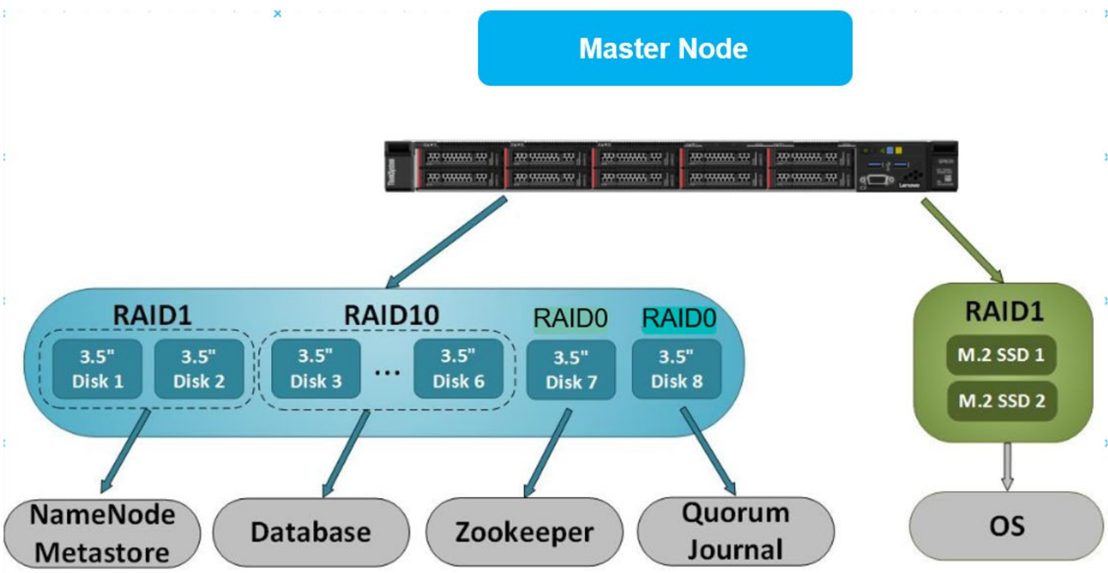

Figure 10. Cloudera Master and Utility node disk assignment

6.2.3 System Management and Gateway Nodes

Known as Gateway, System Management or edge nodes, it is a control point managing user access and connecting the Cloudera cluster to an outside network for remote administration access, for ingesting data from an outside source, or for running end user application software which accesses the Cloudera Data Platform cluster.

A system management/gateway node can start from a minimal node configured for remote administration of the Linux OS and for hardware maintenance. Based on the particular requirements of the cluster for high speed ingesting of data and edge node applications, the CPU, memory, storage, and network capability of this server can be increased.

6.2.4 OpenShift Cluster Nodes for CDP Private Cloud Experiences

Cloudera Data Platform (CDP) Private Cloud Experiences requires hardware for a dedicated OpenShift Container Platform (OCP) cluster. An OpenShift cluster consists of several master nodes for managing OpenShift and many worker nodes for running your application on CDP. CDP Private Cloud Experiences deployment consists of a Private Cloud Management Console and one or more environments that are created for deploying the experiences. The Management Console is a service used by CDP administrators to manage environments, users, and services. The number of worker nodes and worker nodes configurations needed depends on factors such as the number of virtual warehouses or machine learning workspaces required for your workloads.

6.3 Cluster Software Stack

6.3.1 Cloudera Data Platform CDP

The following Cloudera CDP and OpenShift software components were installed for this reference design guide

Table 4. Cloudera CDP and OpenShift software components

Component	Version
CDP Private Cloud	7.1
Operating System	Red Hat Enterprise Linux Server release 8
OpenShift	4.8/4.10

6.3.2 HDFS and Ozone

Hadoop Distributed File System (HDFS) is a Java-based file system for storing large volumes of data. Designed to span large clusters of commodity servers, HDFS provides scalable and reliable data storage.

Apache Hadoop Ozone is a scalable, redundant, and distributed object store optimized for big data workloads. Ozone supports both HDFS API and S3 API. Ozone can be co-located with HDFS with unified security and governance policies and offer seamless application portability. Apart from scaling to billions of objects of varying sizes, Ozone can function effectively in containerized environments such as Kubernetes and YARN.

Main Differences between HDFS and Ozone:

- HDFS supports up to 500 million files, while Ozone can support billions of files
- HDFS supports up to 100 TB/node, while Ozone can support 400 TB/node
- HDFS supports up to 8TB drives, while Ozone supports 16TB drives

6.4 Cloudera Service Role Layouts

Because the Master node is responsible for many memory-intensive tasks, multiple Master and Utility nodes are needed to split out functions. For most implementations, the size of the Cloudera cluster drives how many Master/Utility nodes are needed. Table 5 and Table 6 provide a high-level guideline for a cluster that provides HA NameNode and ResourceManager failover when configured with multiple Master nodes.

Table 5. Service Layout Matrix for High Availability

Node	Master Node 1	Master Node 2	Master Node 3	Worker Nodes	Utility Host 1	Utility Host 2	Gateway Hosts
Service /Roles	NameNode JournalNode	NameNode JournalNode	JournalNode	DataNode	Cloudera Manager		Hue
	ZooKeeper	ZooKeeper	ZooKeeper		Cloudera Manager Management Service		HiveServer2
	YARN Resource-Manager	YARN Resource-Manager		NodeManager	Cruise Control		Gateway configuration
	JobHistory Server			Impalad	Hive Metastore		
	Kudu master	Kudu master	Kudu master	Kudu tablet server	Impala Catalog Server		
	Failover Controller	Failover Controller		Kafka Broker	Impala StateStore		
				Kafka Connect	Oozie		
	HBase master	HBase master		HBase RegionServer	Ranger Admin, Tagsync, Usersync servers	Ranger Admin server	
	Schema Registry	Schema Registry	Spark History Server		Atlas server	Atlas server	

				Solr server	Solr server	Solr server	
					Streams Messaging Manager		
				Streams Replication Manager Driver	Streams Replication Manager Service		

As the number of worker nodes increases in the cluster, more Master and Utility nodes will be needed as shown in Table 6.

Table 6. Cluster size and node types

Number of Worker Nodes	Number of Master Nodes	Number of Utility nodes	Number of Gateway Nodes
3 - 20	3	2	1+
20 - 80	3	2	1+
80 - 200	3	8	1+
200 - 500	5	8	1+

Note: *To ease scale-up of worker nodes, one should plan ahead by installing the next level of Master nodes to be ready as Worker nodes cross one of the boundaries.*

Installing and managing the Cloudera Stack

The Hadoop ecosystem is complex and constantly changing. Cloudera makes it simple so enterprises can focus on results. Cloudera Manager is the easiest way to administer Hadoop in any environment, with advanced features like intelligent defaults and customizable automation. Combined with predictive maintenance included in Cloudera's Support Data Hub, Cloudera Data Platform keeps the business up and running.

Reference Cloudera's latest Installation documentation for detailed instructions on Installation:
https://docs.cloudera.com/cdp-private-cloud-base/latest/installation/topics/cdpdc-installation.html

Reference Cloudera software service layout recommendations:
https://docs.cloudera.com/documentation/enterprise/latest/topics/cm_ig_host_allocations.html#host_role_assignments

6.5 System Management

Systems management of a cluster includes Operating System, Hadoop & Spark applications and hardware management. Systems management uses Cloudera Manager and is adapted from the standard Hadoop distribution, which places the management services on separate servers than the worker servers. The Master node runs important and high-memory use functions, so it is important to configure a powerful and fast server for systems management functions. The recommended Master node hardware configuration can be customized according to client needs.

Hardware management uses the Lenovo XClarity™ Administrator, which is a centralized resource management solution that reduces complexity, speeds up response and enhances the availability of Lenovo server systems and solutions. XClarity™ is used to install the OS onto new worker nodes; update firmware across the cluster nodes, record hardware alerts and report when repair actions are needed.

Figure 11 shows the Lenovo XClarity™ Administrator interface in which servers, storage, switches and other rack components are managed and status is shown on the dashboard. Lenovo XClarity™ Administrator is a virtual appliance that is quickly imported into a server virtualized environment.

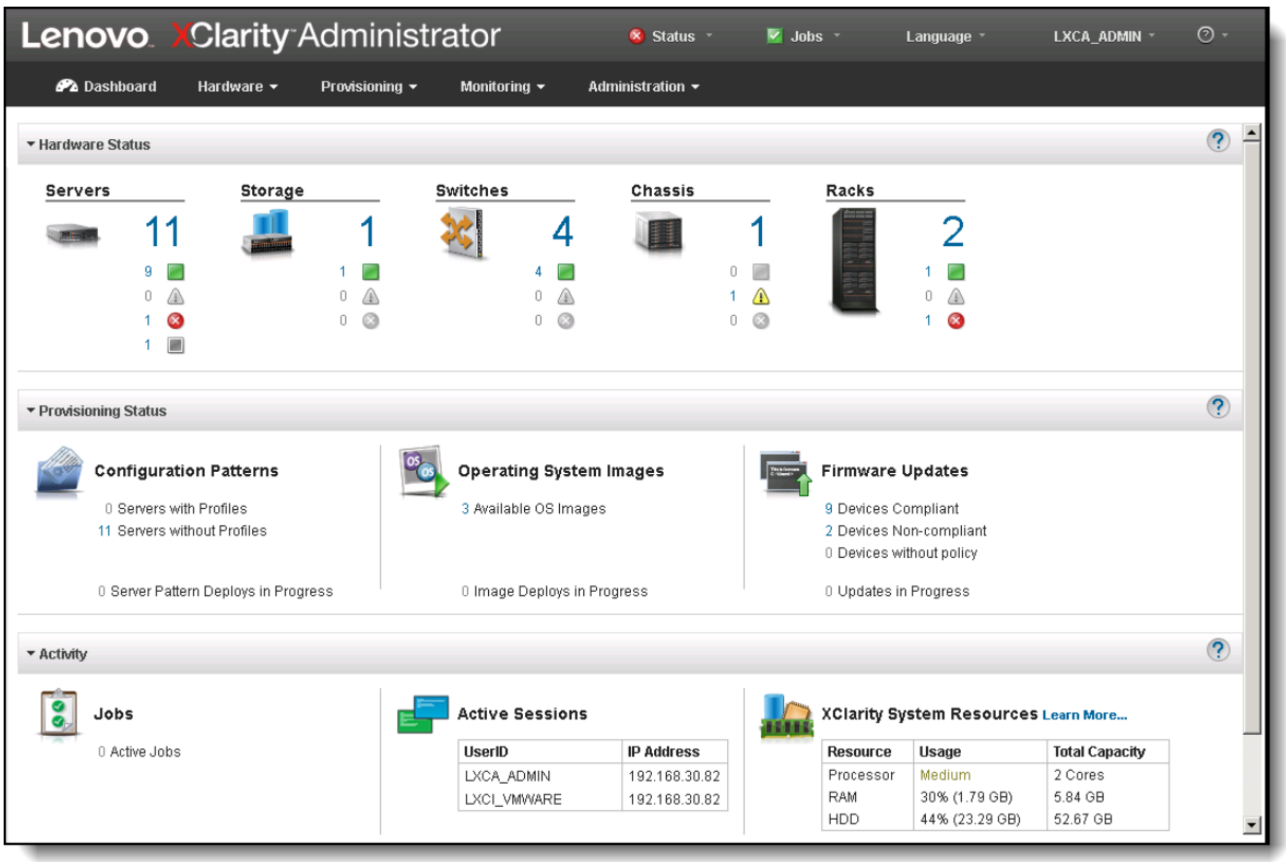

Figure 11. XClarity™ Administrator interface

In addition, the xCAT open source cluster management tool is available which uses a command line interface to provide a unified interface for hardware control, discovery and operating system deployment. It can be used to facilitate or automate the management of small and large cluster sizes. For more information about xCAT, see Resources section on page 34.

6.6 **Networking**

The reference architecture specifies two networks: a high-speed data network and a management network. Two types of rack switches are required: one 1Gb for out-of-band management and a pair of high-speed data switches for the data network with High Availability. See Figure 12 below.

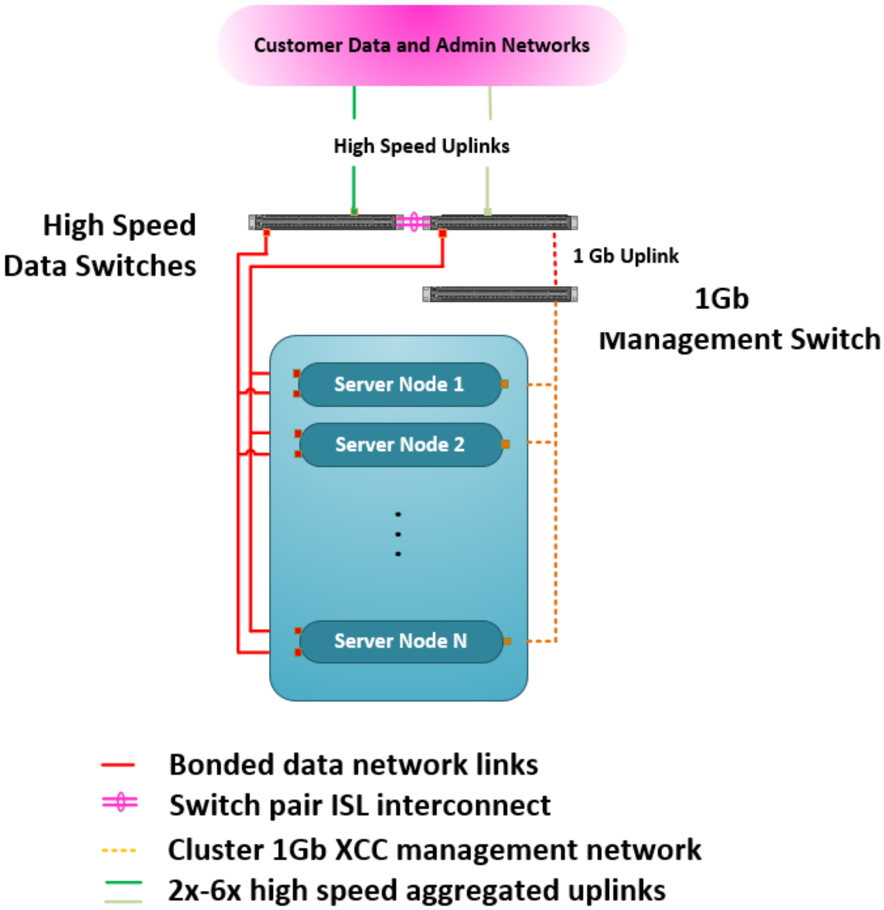

Figure 12. Cloudera network

6.6.1 Data Network

The data network creates a private cluster among multiple nodes and is used for high-speed data transfer across worker and master nodes, and also for importing data into the Cloudera cluster. The Cloudera cluster typically connects to the customer's corporate data network. This reference design architecture demonstrates 10Gb network switches with 40Gb uplink ports. User can use other high speed(25G/100G) network switches.

The two ethernet NIC ports of each node are link aggregated into a single bonded network connection giving up doubling the bandwidth of each individual link. Link redundancy is also provided if one link fails. The two data switches are connected together as a Virtual Link Aggregation Group (vLAG) pair using LACP to provide the switch redundancy. Either high speed data switch can drop out of the network and the other switch continues transferring traffic. The switch pairs are connected with dual 10Gb links called an ISL, which allows maintaining consistency between the two peer switches.

6.6.2 Hardware Management Network

The hardware management network is a 1GbE network for out-of-band hardware management. Through the XClarity™ Controller management module (XCC) within the Lenovo ThinkSystem servers, the out-of-band network enables hardware-level management of cluster nodes, such as node deployment, UEFI firmware configuration, hardware failure status and remote power control of the nodes.

Hadoop has no dependency on the XCC management function. The Cloudera OS/management network can be shared with the XCC hardware management network, or can be separated via VLANs on the respective switches. The Cloudera cluster and hardware management networks are then typically connected directly to the customer's existing administrative network to facilitate remote maintenance of the cluster.

6.6.3 Multi-rack Network

The data network in the predefined reference architecture configuration consists of a single network topology. A rack consists of redundant high-speed data switches with two bonded links to each server node. Additional racks can be added as needed for scale out. Beginning with the third rack a core switch with 40G and 100Gb uplinks for rack aggregation is the best choice for this purpose.

Figure 13 shows a 2-rack configuration. A single rack can be upgraded to this configuration by adding the second rack with the LAG network connection show.

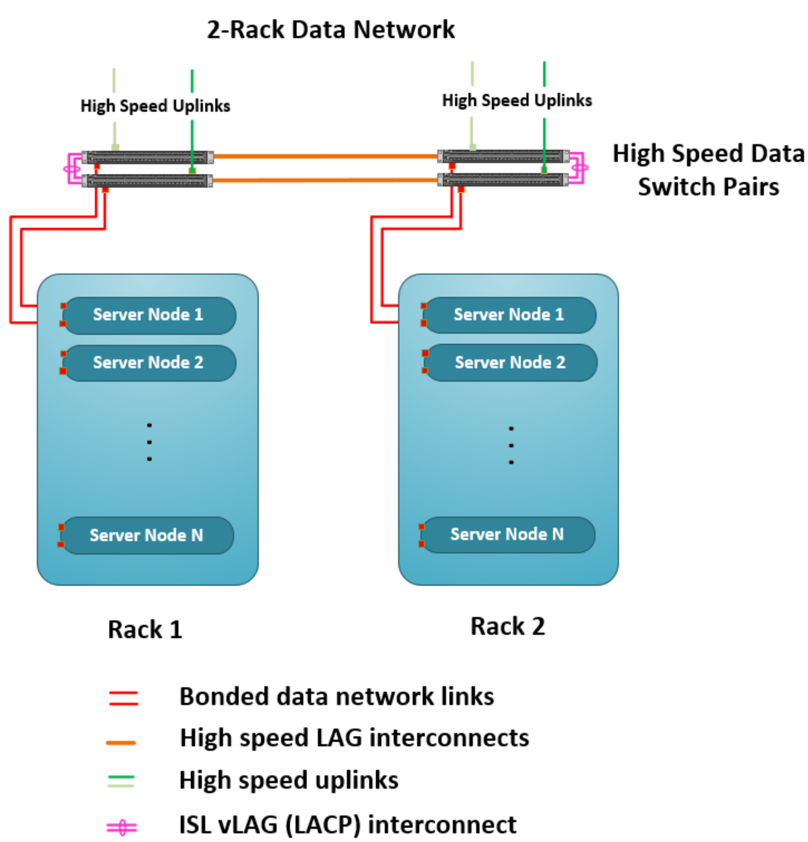

Figure 13. Cloudera 2-rack network configuration

Figure 14 shows how the network is configured when the Cloudera cluster contains 3 or more racks. The data network is connected across racks by four aggregated uplinks from each rack's high speed data switch to a core switch. The 2-rack configuration can be upgraded to the 3-rack configuration as shown. Additional racks can be added with similar uplink connections to the cross rack switch.

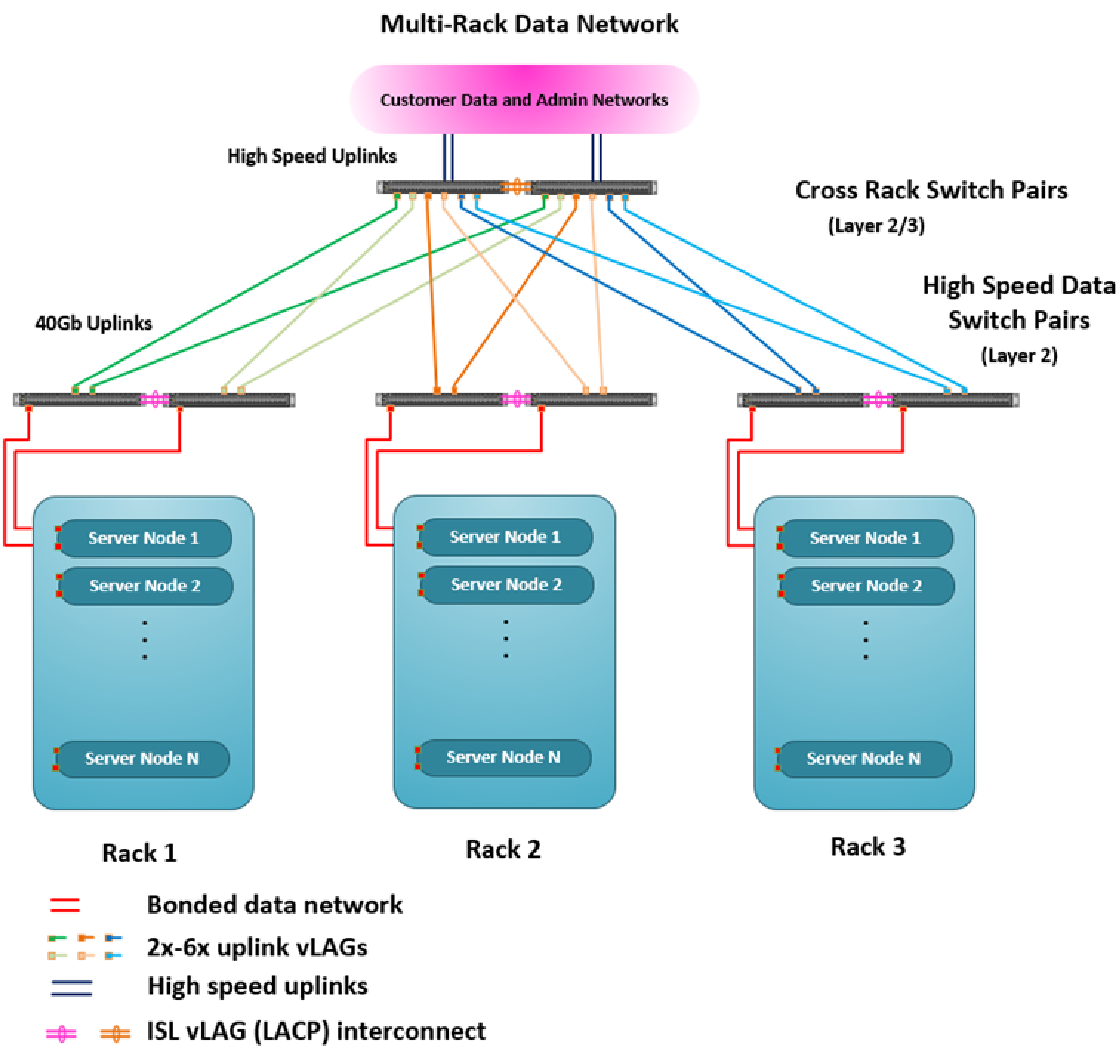

Figure 14. Cloudera multi-rack rack network configuration

Within each rack, 1Gb management switch can be configured to have two uplinks to the data switches for propagating the management VLAN across cluster racks through the cross-rack switch. Other cross rack network configurations are possible and may be required to meet the needs of specific deployments and to address clusters larger than three racks.

For multi-rack solutions, the Master nodes can be distributed across racks to maximize fault tolerance.

6.7 **Predefined Cluster Configurations**

The intent of the predefined configurations is to aid initial sizing for customers and to show example starting points for four production cluster sizes: starter rack, half rack, full rack, and a 3 rack multi-rack configuration. These consist of Worker nodes, Master/Utility nodes, system management/gateway nodes, network switches, storage enclosures and rack hardware. Figure 15 show rack diagrams with a description of each component.

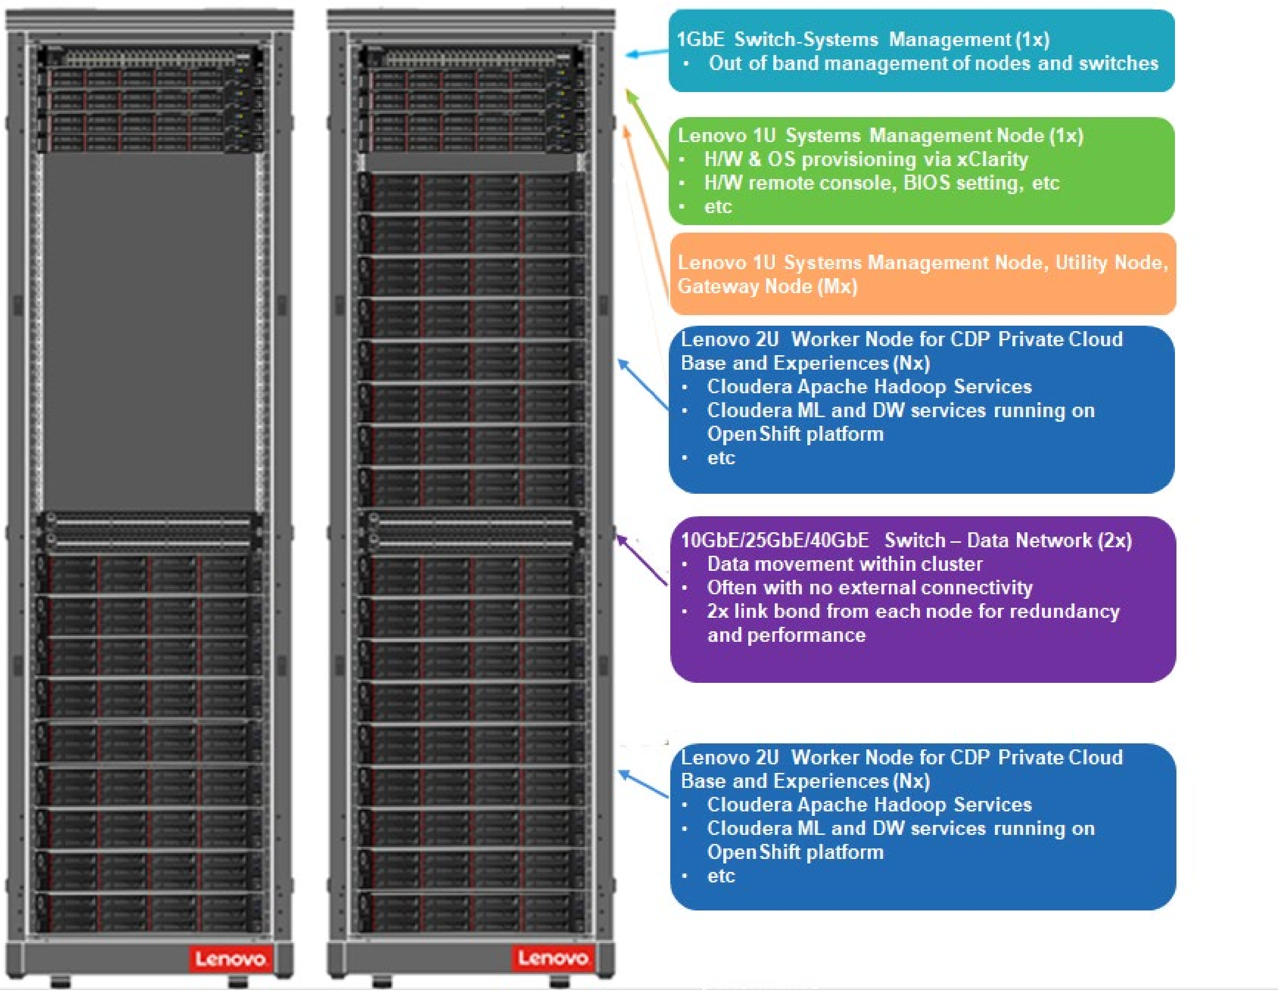

Half Rack **Full Rack**

Figure 15. Half rack and full rack Cloudera predefined configurations

7 Customer Case

Lenovo and Cloudera empower customers in various industries. We have supported people in Banking, Insurance, Retail, Telecom, Automotive, and many more industries. We together can build a smart data platform with stable, reliable, and powerful Lenovo ThinkSystem servers.

One bank customer has been providing its customers with personalized services to meet their financial needs for several years. However, with the growing complexity of financial services and the increasing competition in the market, the Bank realized that it needed to modernize its IT infrastructure and data analytics capabilities to enhance customer experience.

After conducting extensive research, the Bank chose Cloudera CDP as its data analytics platform, and Lenovo ThinkSystem servers as the infrastructure to power it. The Cloudera CDP platform is known for its ability to handle large amounts of data, provide advanced data analytics and machine learning capabilities, and ensure data security and governance. Lenovo ThinkSystem servers, on the other hand, are known for their stability, reliability, and power, making them an ideal choice to run the Cloudera CDP platform.

We worked together with people in bank, and designed a Cloudera CDP platform with suitable master node and work nodes configurations to support short term and middle term services. And the design can be extended and upgrade for long term requirements.

With the data analytics platform in place, the Bank was able to collect, store, process, and analyze large amounts of customer data in real-time and in batch. This enabled the bank to gain insights into customer needs and preferences, which it used to provide personalized financial solutions that meet the unique needs of each customer.

With the data analytics platform deployed on stable and reliable Lenovo ThinkSystem servers, The Bank was able to provide stable and reliable services to improve its customer experience significantly. They led to increased customer satisfaction and loyalty.

The Bank can leverage Cloudera CDP and Lenovo ThinkSystem servers to improve its customer experience significantly. With the new data analytics capabilities, The Bank is well-positioned to remain competitive in the industry and continue to provide its customers with the best possible services.

8 Resources

For more information, see the following resources:

Data: A Lenovo Solutions Perspective

- Lenovo Press whitepaper: https://lenovopress.com/lp1367-lenovo-solutions-perspective-on-data

Lenovo ThinkSystem SR645 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1607-thinksystem-sr645-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1627-sr645-v3-3d-tour

Lenovo ThinkSystem SR665 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1608-thinksystem-sr665-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1628-sr665-v3-3d-tour

Lenovo ThinkSystem SR655 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1610-thinksystem-sr655-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1630

Lenovo ThinkSystem SR635 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1609-thinksystem-sr635-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1629

Lenovo ThinkSystem SR650 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1601-thinksystem-sr650-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1621-thinksystem-sr650-v3-interactive-3d-tour

Lenovo ThinkSystem SR630 V3 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1600-thinksystem-sr630-v3-server
- 3D Tour: https://lenovopress.lenovo.com/lp1620-thinksystem-sr630-v3-interactive-3d-tour

Lenovo ThinkSystem SR645 server:

- Lenovo Press product guide: https://lenovopress.com/lp1280-thinksystem-sr645-server
- 3D Tour: https://www.lenovofiles.com/3dtours/products/superblaze/sr645/index.html

Lenovo ThinkSystem SR665 server:

- Lenovo Press product guide: https://lenovopress.com/lp1269-thinksystem-sr665-server
- 3D Tour: https://www.lenovofiles.com/3dtours/products/superblaze/sr665/index.html

Lenovo ThinkSystem SR650 V2 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1392-thinksystem-sr650-v2-server
- 3D Tour: https://lenovopress.lenovo.com/lp1424-lenovo-thinksystem-sr650-v2-3d-tour

Lenovo ThinkSystem SR630 V2 server:

- Lenovo Press product guide: https://lenovopress.lenovo.com/lp1391-thinksystem-sr630-v2-server
- 3D Tour: https://lenovopress.lenovo.com/lp1423-3d-tour-lenovo-thinksystem-sr630-v2

AMD EPYC™ processor:

- Product page: https://www.amd.com/en/products/epyc-server

Intel Xeon Scalable processor:

- Product page: https://www.intel.com/content/www/us/en/products/processors/xeon.html

Lenovo XClarity Administrator:

- Product page: https://lenovopress.com/tips1200-lenovo-xclarity-administrator
- Lenovo Press product guide: https://lenovopress.com/tips1200.pdf

Lenovo ThinkSystem GPU options:

- GPU options: https://lenovopress.com/lp0767-gpu-options-for-thinksystem-servers

Cloudera:

- Cloudera Data Platform (CDP): https://www.cloudera.com/products/open-source/apache-hadoop/key-cdh-components.html
- Cloudera Installation Guide: https://docs.cloudera.com/cdp-private-cloud-base/latest/installation/topics/cdpdc-installation.html
- Cloudera products and services: https://www.cloudera.com/products.html
- Cloudera solutions: http://www.cloudera.com/content/cloudera/en/solutions.html
- Cloudera resources: https://www.cloudera.com/resources.html

Red Hat: https://www.redhat.com/en

Red Hat OpenShift: https://www.openshift.com/

NVIDIA:

- Accelerating Apache Spark 3.X : https://www.nvidia.com/en-us/deep-learning-ai/solutions/data-science/apache-spark-3/ebook-sign-up/
- Spark-Rapids on-prem: https://nvidia.github.io/spark-rapids/docs/get-started/getting-started-on-prem.html
- NVIDIA A100: https://www.nvidia.com/en-us/data-center/a100/

Open source software:

- Hadoop: hadoop.apache.org
- Spark: spark.apache.org
- Solr: solr.apache.org
- Flume: flume.apache.org
- HBase: hbase.apache.org
- Hive: hive.apache.org
- Hue: gethue.com
- Impala: rideimpala.com

- Oozie: oozie.apache.org
- Mahout: mahout.apache.org
- Pig: pig.apache.org
- Sentry: entry.incubator.apache.org
- Sqoop: sqoop.apache.org
- Whirr: whirr.apache.org
- ZooKeeper: zookeeper.apache.org
- Parquet: parquet.apache.org
- Arrow: arrow.apache.org
- NiFi: https://nifi.apache.org/
- Hadoop Virtualization Extensions (HVE): https://issues.apache.org/jira/browse/HADOOP-8468
- xCat: xcat.org

Document history

Version 1.0	March 2021	First version
Version 2.0	June 2021	Updated Server Description section
Version 3.0	September 2021	Updated Apache Spark 3.0 with GPU acceleration
Version 4.0	April 2023	Add Intel V3 server description, and Use case. Update software version.
Version 4.1	June 2023	Add AMD V3 server description.
Version 4.2	March 2024	Add Intel Emerald Rapids processors description. Update GPU section.
Version 4.3	Novmember 2024	Updated with SR665/SR645 V3 with AMD Turin processors
Version 4.4	June 2025	Updated with SR650/SR630 V4 servers

Trademarks and special notices

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.