Unlock the Power of Private Cloud Big Data Analytics

Transform complex data into clear and actionable insights with a high-performance, private cloud big data analytics solution from Cloudera and Intel

Is this solution for you? Do you...

- ✓ Need more performance from your Cloudera distribution?
- ✓ Need to integrate isolated data silos into a cohesive data lake?
- ✓ Have large amounts of on-premises data?
- ✓ Want to reduce total costs by replacing expensive proprietary data warehouses?
- ✓ Need a platform for easy and fast distributed data analytics and data management?
- ✓ Need to consolidate hardware sprawl?
- ✓ Need to increase data center efficiency and flexibility?

Learn about the Business Story →
Learn more about the Reference Architecture →

Intel Authors

Esther Baldwin
Cloud & Enterprise Solutions Lead

Angela Gill
Cloudera Alliance Manager

Swepna Mohan-Doss
Cloud Solution Engineer

Amandeep Raina
Cloud Solution Engineer

Lifeng Wang
Machine Learning Engineer

Ning Wang
Machine Learning Engineer

June Wong
Software Program Manager

Cloudera Author

Ali Bajwa
Partner Engineering Director

Contributors

A Xida
Software Engineer

Cheng Xu
Software Engineering Manager

Ye Yuqiang
Senior Software Engineer

Business Story
Unlock the Power of Private Cloud Big Data Analytics

Solution Benefits

- Open-source platform provides flexibility and interoperability with existing tools.
- Disaggregation of compute and storage leads to more flexibility and efficiency.
- A comprehensive set of management tools simplifies cluster configuration and scaling.
- Upgrading from a legacy distribution to CDP Private Cloud Base running on 2nd Generation Intel® Xeon® Scalable processors provides up to a 2x performance improvement.[1]
- Using Intel® Optane™ PMem for Apache Kudu block cache can provide a 6.3x increase in throughput and a 13.38x decrease in latency, compared to a DRAM-based configuration.[2]

Executive Summary

It is well understood that enterprises can extract business value from the large volume of data they generate. The difficulty lies in integrating isolated silos of data throughout the business and managing data efficiently. One common hurdle is competing business units that control access to the data but refuse to share it. Also, internal policies that were intended to protect the company and its customers can hamper data analysis and create a significant burden to extracting business value and improving decision making. The velocity of the data flood adds additional stress on business units, IT, and executive management, who must work through the complexities created in a data-driven world. Every enterprise has data management problems. The difficulty is identifying and then correcting them promptly. Simply having a better product or service is insufficient in today's world. Getting your product or service to the right customer, whether new or existing, is paramount. Spotting trends before your biggest rival could mean the difference between success or failure.

If these challenges sound familiar, you're not alone. According to Gartner, 91 percent of organizations struggle to reach data maturity.[3] But don't despair—a collaboration between Intel and Cloudera has created a big data analytics platform specifically designed for large-scale on-premises workloads.

Cloudera Data Platform (CDP) Private Cloud powers on-premises, data-driven decision making by easily, quickly, and safely connecting and securing the business's entire data lifecycle. This big data analytics platform helps business leaders modernize their data center by streamlining data management and workload orchestration. Separation of compute and storage leads to improved flexibility and efficiency. With CDP Private Cloud, enterprises can migrate to a container-based environment and take advantage of the agility and scalability of containers.

Once data volume, velocity, and data access issues are resolved, time to value becomes the issue. Solving a problem one minute too late is lost business value. Said another way, "time is money." With the Intel® architecture underlying CDP Private Cloud, we provide the power that big data analytics demand. Intel and Cloudera collaborated to improve compute performance, storage efficiency, artificial intelligence (AI) acceleration, and more. The result is a private cloud data platform built to meet today's big data analytics needs that can scale to meet your business needs today and into the future. Tests show that a modern version of CDP Private Cloud running on the latest Intel® hardware can improve the data analysis performance of several aspects of the CDP Private Cloud system. For example, upgrading to CDP Private Cloud from a legacy distribution, **running on 2nd Generation Intel® Xeon® Scalable processors, can improve throughput by up to 2.23x.**[4] Using Intel® Optane™ persistent memory (PMem) for Apache Kudu block cache can **provide up to a 6.3x increase in throughput and a 13.38x decrease in latency,** compared to an all-DRAM-based configuration (depending on dataset size).[5]

This document provides a business-level overview of CDP Private Cloud, describes a reference solution for deployment, and highlights the platform's performance and scalability. And if you are already using a legacy distribution of a Cloudera product, this document also describes best practices for migrating to the latest distribution.

Business Challenge: Extracting Business Insights from a Jumble of Data

Data is everywhere in your enterprise. It is constantly being generated by machines, customers, and applications. It piles up in various data warehouses. The significant problem is that hidden in that data are insights about your business—information about network security, customer preferences, supply chain dynamics, and more. When data exists in silos and data analytics runs in those same silos, it is nearly impossible to unlock the business value that lies within this volume of data. And as data continues to grow and big data analytics workloads increase accordingly, your infrastructure must be able to scale appropriately. However, storage requirements often outpace compute needs so the ability to scale these resources independently is crucial to data center efficiency.

CDP Private Cloud is an open source, scalable data platform optimized to run on high-performance Intel® hardware. It also supports the disaggregation of compute and storage. Using CDP Private Cloud, you can quickly and cost-efficiently extract value from your data without leaving the data center (see the Solution Value section).

Use Cases Abound for a Private Cloud Data Analytics Platform

Many industries, from manufacturing to healthcare to retail, and from transportation to hospitality to life sciences, are under pressure to turn their data into business value. CDP Private Cloud can be used across several broad use cases, including data lakes; extract, load, and transform (ELT) applications; and offloading analytics from expensive proprietary databases. And while the trend is to take data to the cloud (and Cloudera does support cloud deployments), keeping data on-premises in a private cloud sometimes makes more sense. For example, the data may be sensitive (such as intellectual property). Or the analytics use case may require extremely low latency (such as real-time fraud detection).

The following examples illustrate how you can put CDP Private Cloud to work to solve your business's data processing issues. These use cases demonstrate the flexibility of the CDP Private Cloud solution:

- **Manufacturing.** Send Internet of Things (IoT) data—much of it semi-structured—into a data lake, then run analytics on that data for predictive maintenance, real-time production line changes, and more. Intel manufacturing uses large datasets containing billions of data points per day per factory. For high-volume manufacturing companies like Intel, advanced analytics using the latest Intel® processors and CDP Private Cloud can help reduce data processing time and can lead to improved yield, fast time to market, accelerated insights, increased productivity, excursion control, and more.
- **Healthcare and life sciences.** Use a data lake to store unstructured data (phone call recordings, webinar transcripts, imaging data, etc.) and make that data available to academic researchers. IQVIA, a global provider of advanced analytics, technology solutions, and contract research services to the life sciences industry, has used CDP Private Cloud to accelerate query responses from days to seconds.[6]
- **Sales and marketing.** Use data to optimize marketing campaigns and create advanced recommendation engines.
- **Financial services.** Use extract/transform/load to detect fraud, predict customer churn, or perform risk management.

About Cloudera

CLOUDERA

Cloudera, Inc. is a U.S.-based software company that provides a software platform for data engineering, data warehousing, machine learning, and analytics that runs in the cloud or on-premises. Founded in 2008, the company is committed to accelerating enterprise-class data management innovation to make what is impossible today, possible tomorrow. Cloudera solutions empower enterprises to transform complex data into clear and actionable insights.

Powered by the relentless innovation of the open source community, Cloudera has offices around the globe and is headquartered in Palo Alto, California.

- **Supply chain management.** Use data science to identify cost savings opportunities, speed the costing cycle, perform historical pricing analysis, and conduct "what if" analysis for procurement and planning.
- **Transportation.** Gather data from sensors and run real-time analytics that provide input to autonomous cars, or analyze images for surveillance and safety systems.
- **Cybersecurity.** Pour all networking data into a data lake and run analytics to detect vulnerabilities and threats.

> "...upgrading to CDP Private Cloud from a legacy distribution, running on 2nd Gen Intel Xeon Scalable processors, can improve throughput by up to 2.23x.[1]"

Solution Value: High Performance, Ultimate Flexibility, and Excellent Scalability

Enterprises need answers—and they need them now! At its core, CDP Private Cloud is a next-generation, cloud-native architecture for on-premises deployments that drives real-time and batch analytics processing. It speeds time to value by separating compute and storage, integrating a suite of security and governance tools, and packaging all pieces together with a simpler and intuitive management console.

CDP Private Cloud delivers a suite of analytic engines ranging from stream and batch data processing to data warehousing, operational database, and machine learning (see Figure 1). By using CDP Private Cloud, enterprises gain end-to-end big data capabilities—ingest, store, process, and analyze for insights. The CDP Private Cloud shared data experience applies consistent security and governance, enabling users to share and discover data for use across many demanding big data analytics workloads.

Manage and Secure the Data Lifecycle with Cloudera

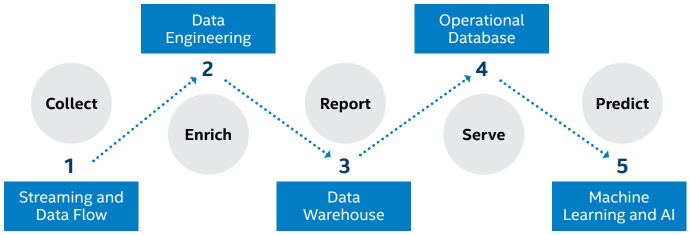

Figure 1. CDP Private Cloud simplifies data management, offers a wide variety of analytic engines, and provides a shared data experience.

Some specific business benefits from CDP Private Cloud include:

- **Cloudera.** An end-to-end data analytics processing platform, in contrast to standalone solutions that require stitching together for the data flow.
- **Ultimate agility and flexibility.** It's easy to integrate CDP Private Cloud with existing infrastructure and tools, and the platform offers robust security, governance, data protection, and data management features. CDP Private Cloud is open source, which allows for flexibility in choosing the technologies you want to use without vendor lock-in. What's more, the open-source community drives the evolution of data management and advanced analytics. Intel and Cloudera's strong relationships with a broad portfolio of data center solution providers can help streamline the process for building solutions.
- **Excellent scalability.** The CDP Private Cloud reference solution from Intel is the foundation for a highly scalable big data analytics platform that can store any amount or type of data in its original form—and keep it for as long as it's needed. Also, CDP Private Cloud includes software that simplifies scaling by automating node configuration.
- **High efficiency.** Support for independent scaling of compute and storage resources lets data centers invest in precisely the resources they need for their specific workloads. Avoiding overprovisioning of either compute or storage can lead to cost savings, a simpler infrastructure, and lower maintenance efforts.
- **Enhanced security and governance.** CDP Private Cloud comes with a complete suite of security and governance capabilities. These services regulate what end users can do through the analytic experiences but operate independently of these experiences. This means the security and governance tools can be independently configured, managed, and upgraded, and these changes will automatically get reflected in the analytic experiences.

Intel and Cloudera's participation in joint engineering results in optimizations of CDP Private Cloud pertaining to faster compute performance, increased storage efficiency, enhanced security, excellent AI support, and great query performance. These optimizations let CDP Private Cloud take advantage of Intel® Optane™ PMem, Intel Xeon Scalable processors, Intel® FPGAs, Intel® QuickAssist Technology, Intel® Advanced Vector Extensions 512, Intel® Math Kernel Library, Intel® AES-NI, and Intel® Intelligent Storage Acceleration Library. **Intel testing shows that running CDP Private Cloud on 2nd Generation Intel Xeon Scalable processors can boost throughput by up to 2.23x.**[7] For workloads using Apache Kudu, using Intel Optane PMem for Kudu block cache **can increase throughput by up to 6.3x and decrease latency by up to 13.38x** (compared to an all-DRAM-based configuration and depending on dataset size).[8]

Additional benefits of CDP Private Cloud include **manageability features,** including native high-availability; fault-tolerance and self-healing storage; automated backup and disaster recovery; advanced system and data management; and a single pane of glass for cluster administration, automation, management, and security.

Solution Architecture: Scalable and Agile Big Data Analytics Platform

CDP Private Cloud (see Figure 2) is built on and optimized for Intel® compute, storage, and networking architecture. The reference solution for CDP Private Cloud incorporates the following Intel technologies:

- **2nd Generation Intel Xeon Scalable processors.** These processors are optimized for big data analytics workloads like Hadoop. They incorporate architecture improvements and enhancements for compute-intensive and data-intensive workloads, making them well suited for ingesting and analyzing massive quantities of data.
- **Intel® 3D NAND SSDs.** These PCIe/NVMe-based SSDs deliver scalable, cost-effective performance and low latency. The SSDs also offer outstanding quality, reliability, advanced manageability, and serviceability to minimize service disruptions.
- **Intel® Ethernet network connection.** Intel Ethernet network controllers, adapters, and accessories enable agility in the data center to deliver services efficiently and cost effectively. Compatible with the Open Compute Platform, these high-performance connectors support high throughput, reliability, and compatibility.
- **Intel® Optane™ technology.** The first breakthrough in memory and storage in 25 years, Intel Optane PMem and Intel® Optane™ SSDs are unique innovations that bridge critical gaps in the storage and memory hierarchy, delivering persistent memory, large memory pools, fast caching, and fast storage.

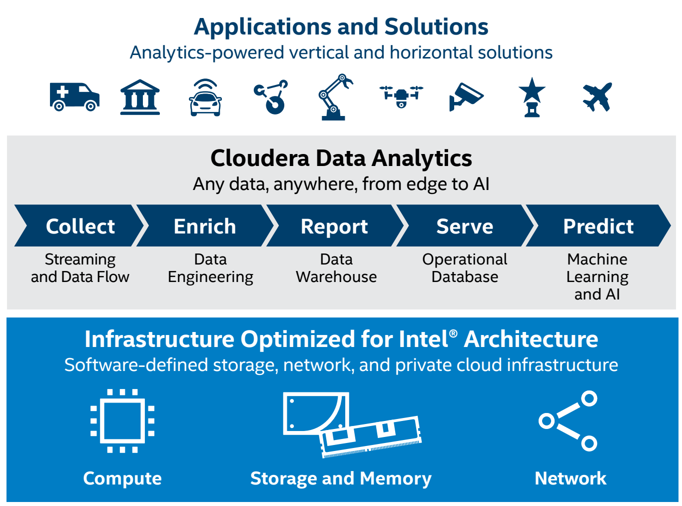

Figure 2. High-level diagram of the basic Cloudera Data Platform Private Cloud solution architecture.

Ready to learn more? Turn the page for a detailed Reference Architecture discussion.

Reference Architecture
Unlock the Power of Private Cloud Big Data Analytics

Get better performance by upgrading to a newer version of Cloudera Data Platform with 2nd Gen Intel® Xeon® Scalable processors.[9] Using Intel® Optane™ PMem for Apache Kudu block cache can increase throughput and decrease latency.[10]

Big Data Analytics Process

1. **Ingest.** This reference solution supports high-volume data ingestion of structured and unstructured data from various sources such as transactional relational database management systems, operations data, web logs, click streams, and other external sources.
2. **Prepare.** Once ingested, the data is cleaned and formatted with metadata and schema.
3. **Analyze.** Next, the data is loaded into the analytical data warehouse with shared local storage applied with pre-defined use-case logic for sorting and combining functions on distributed computing nodes.
4. **Act.** Finally, results in the form of compressed datasets are made available for business consumption to run reporting, machine-learning models, and business intelligence. In addition, the solution is flexible enough to support ad-hoc analysis of data when there are no pre-defined use cases.

Table of Contents

Overview of CDP Private Cloud

CDP Private Cloud provides a scalable, versatile, and integrated platform that helps modernize the data center. This platform simplifies managing the growing volume and variety of data in your enterprise, thereby unleashing the business value of that data. By disaggregating compute and storage and supporting a container-based environment, CDP Private Cloud helps enhance business agility and flexibility. Using Cloudera products and solutions can help you deploy and manage Apache Hadoop and related projects, manipulate and analyze your data, and enhance data security.[11]

CDP Private Cloud includes two versions: CDP Private Cloud Base and CDP Private Cloud Experiences. The Base edition includes shared data experience, storage management, and traditional bare-metal data lifecycle analytics. It replaces—and is equivalent to—CDP Data Center. It is the foundation of CDP Private Cloud. CDP Private Cloud Experiences has access to all the Base features and adds container-based analytic experiences for data warehousing and machine learning. It also adds container-based management and control plane services. As shown in Figure 3, CDP Private Cloud and CDP Public Cloud can be managed from a single pane of glass—the CDP Management Console.

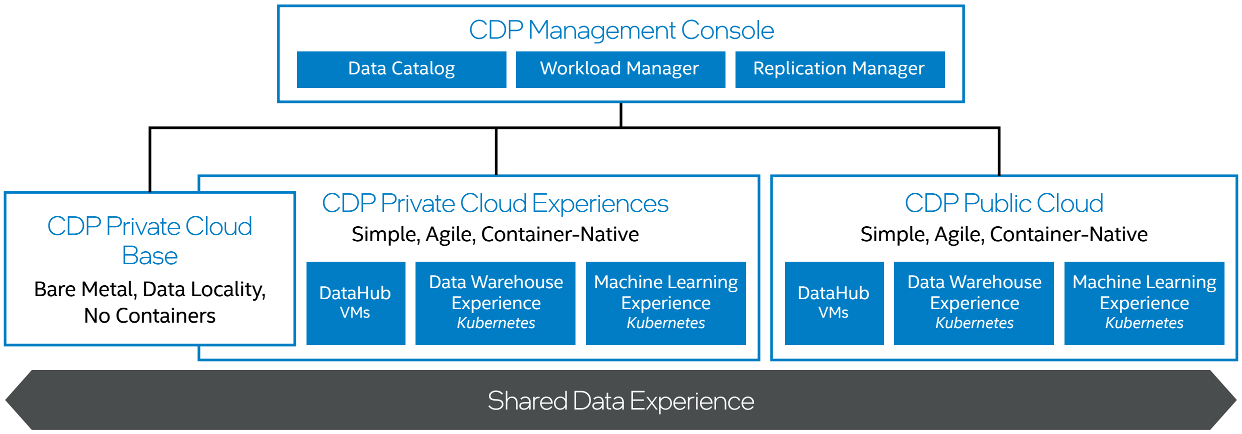

Figure 3. Overview of the Cloudera Data Platform.

Infrastructure Guidance

CDP Private Cloud infrastructure has three primary components:

- **CDP Private Cloud Base.** This houses Hadoop Distributed File System (HDFS) data, Hive metastore (HMS), and Ozone, and it serves as the storage platform for this solution. Apache Ranger and Apache Atlas are also included in CDP Private Cloud Base for data auditing and governance.
- **Red Hat OpenShift.** The minimum requirements for this infrastructure are provided in the CDP Private Cloud Experiences Worker Nodes section.
- **Appropriate networking infrastructure.** There must be sufficient bandwidth between the storage and compute environments. The networking requirements are provided in the Network Considerations section.

Table 1 summarizes software services that each node provides.

Table 1. CDP Private Cloud Base and CDP Private Cloud Experiences Nodes and Roles

CDP Edition	Physical Node	Software Function
Private Cloud Base	Edge Node	• Hadoop Clients • Cloudera Manager
	Master Nodes 1-4	• NameNode • Resource Manager • ZooKeeper
	Worker Nodes	• DataNode • NodeManager • YARN workloads
Private Cloud Experiences	Master Nodes 1-4	• OpenShift services • Kubernetes services
	Worker Nodes	• Kubernetes operators • Workload pods

CDP Private Cloud Base Master Nodes

At least three master nodes are required. More nodes with the same configuration can be used for edge nodes. This configuration is sized for approximately 1 PB of cluster storage or 250 worker nodes. Table 2 lists the recommended hardware for the CDP Private Cloud Base master nodes.

Table 2. CDP Private Cloud Base Master Node Configuration

Recommended Component	
Processor	2x Intel® Xeon® Gold 6226R processor (16 cores, 2.9 GHz, 22 MB cache)
Memory	192 GB (12x 16 GB 2933 MT/s)
Network	Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)
Storage (HDD)	4x 4 TB 7.2 K RPM SATA 6 Gbps
Storage (SSD)	2x 480 GB Intel® SSD DC S4500 (SAS 12 Gbps)

CDP Private Cloud Base Worker Nodes

A minimum of three worker nodes is required, but we recommend at least five to support high availability and downtime for maintenance. Table 3 lists the recommended hardware for general-purpose CDP Private Cloud Base worker nodes. This processor and memory configuration is ideal for nodes that are also running other services, such as Cloudera Data Warehouse, Cloudera Machine Learning, and Cloudera Data Hub.

Table 3. CDP Private Cloud Base General-Purpose Worker Node Configuration

Recommended Component	
Processor	2x Intel® Xeon® Gold 6248R processor (24 cores, 3.0 GHz, 35.75 MB cache)
Memory	384 GB (12 x 32 GB 2933 MT/s)
Network	Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)
Storage (HDD)	12x 4 TB 7.2 K RPM SATA 6 Gbps
Storage (fast drive)	2x Intel® SSD D7-P5510 NVMe
Storage (SSD)	2x 480 GB SSD SAS mixed-use 12 Gbps

In contrast, if most workloads will be running on the CDP Private Cloud Experiences cluster and the CDP Private Cloud Base cluster will be providing only HDFS storage, then the storage-only worker node configuration shown in Table 4 may be more appropriate.

Table 4. CDP Private Cloud Base Data or Storage-Only Worker Node Configuration

Recommended Component	
Processor	2x Intel® Xeon® Gold 6226R processor (16 cores, 2.9 GHz, 22 MB cache)
Memory	192 GB (12 x 16 GB, 2933 MT/s)
Network	Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)
Storage (HDD)	16x 4 TB 7.2 K RPM SATA 6 Gbps
Storage (SSD)	2x 480 GB Intel® SSD DC S4500 (SAS 12 Gbps)

Note: We recommend Intel® Xeon® Gold processors for best performance when using erasure coding with HDFS. The HDFS erasure coding feature uses the Intel® Storage Acceleration Library, which uses the Intel® AES-NI, SSE, Intel® AVX, Intel® AVX2, and Intel® AVX-512 instruction sets that Intel Xeon Gold processors support.

A slightly different configuration is recommended for CDP Private Cloud Base worker nodes that are running memory-intensive workloads (see Table 5). This configuration may be suitable for workloads that benefit from keeping large datasets in memory. The Intel® Optane™ PMem that is used in the large-memory configuration enables large memory capacity.

Table 5. CDP Private Cloud Base Worker Node Configuration for Memory-Intensive Workloads

Recommended Component	
Processor	2-socket Intel® Xeon® Gold 6248R processor (24 cores, 3.0 GHz, 35.75 MB cache)
Memory (DRAM)	384 GB (12x 32 GB, 2933 MT/s, Dual Rank)
Persistent Memory	12x 128 GB Intel® Optane™ PMem
Network	Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)
Storage (data)	1x 6.4 TB Intel® SSD DC P4610 (2.5in, U.2)

CDP Private Cloud Experiences Master Nodes

The CDP Private Cloud Experiences requires a dedicated OpenShift cluster. The OpenShift cluster consists of a number of master nodes (for managing OpenShift) and a number of worker nodes (for running CDP applications). Table 6 specifies the hardware requirements for the CDP Private Cloud Experiences master nodes.

Table 6. CDP Private Cloud Experiences Master Node Configuration

Recommended Component	
Processor	2x Intel® Xeon® Gold 6226R processor (16 cores, 2.9 GHz, 22 MB cache)
Memory	192 GB (12x 16 GB 2933 MT/s)
Network	Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)
Storage (OS and data)	1x 480 GB Intel® SSD DC S4500 (SAS 12 Gbps)

The following master nodes are required: three OpenShift master nodes; one bootstrap node (it can be converted into an OpenShift worker node after initial deployment); and one Cluster System Admin Host node. This master node configuration is adequate for OpenShift container clusters up to 250 nodes and rarely needs to be customized.

CDP Private Cloud Experiences Worker Nodes

A minimum of four worker nodes are required. For a new deployment, we recommend 10 to 20 worker nodes. OpenShift supports heterogeneous node configurations, and it is possible to create specialized node configurations. We provide two recommended configurations for CDP Private Cloud Experiences worker nodes: typical CDP cloud workloads and memory-intensive CDP workloads.

Typical CDP Cloud Workloads

The recommendations in Table 7 provide a good balance of compute and memory for typical CDP cloud workloads. The configuration assumes the Cloudera Data Warehouse and Cloudera Machine Learning experiences are running on the container cloud while the primary big data storage is provided by HDFS on the CDP Private Cloud Base cluster. Therefore, local storage is limited to just enough to support temporary files, disk caches, and storage for other applications. Contact your Intel account representative for assistance in sizing and customizing any of these configurations.

Table 7. CDP Private Cloud Experiences Worker Node Configuration for Typical Cloud Workloads

Recommended Component	
Processor	2x Intel® Xeon® Gold 6248R processor (24 cores, 3.0 GHz, 35.75 M cache)
Memory	768 GB (12x 32 GB 2933 MT/s, Dual Rank)
Network	Intel® Ethernet Network Adapter E810 CQDA2
Storage (OS and data)	1x 6.4 TB Intel® SSD DC P5510

Network Considerations

At a high level, the network requirements for CDP Private Cloud are similar to those of a Cloudera Manager Virtual Private Cloud (CM VPC) deployment because the model of deployment in both cases is similar—a base cluster that houses HDFS storage and remote compute-only clusters that can read from and write to the base HDFS cluster. But for CDP Private Cloud, the network bandwidth requirements are less stringent than those of the CM VPC because data caching technology is introduced at the compute layer, which is not available in CM VPCs.

This implies that the initial load of data from the remote storage would require significant bandwidth between the compute and storage clusters, subject to the quantity of data being ingested. However, subsequent network bandwidth requirements should be significantly lower.

As shown in Figure 4, a minimum of 10 Gbps (recommended 25 Gbps) guaranteed bandwidth is required between each OpenShift worker node and each CDP Private Cloud Base Data Node. To test the worst case bandwidth, the network should be fully stressed with all OpenShift nodes trying to read or write simultaneously from the CDP Private Cloud Base nodes.

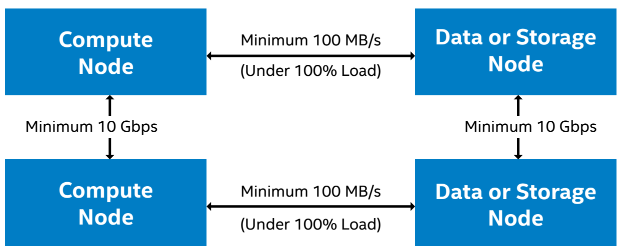

Figure 4. Network bandwidth recommendations.

The recommended network architecture is Spine-Leaf (Figure 5) with no more than a 4:1 oversubscription between the spine and leaf switches. For more information, visit Cloudera's networking documentation for CDP Private Cloud.

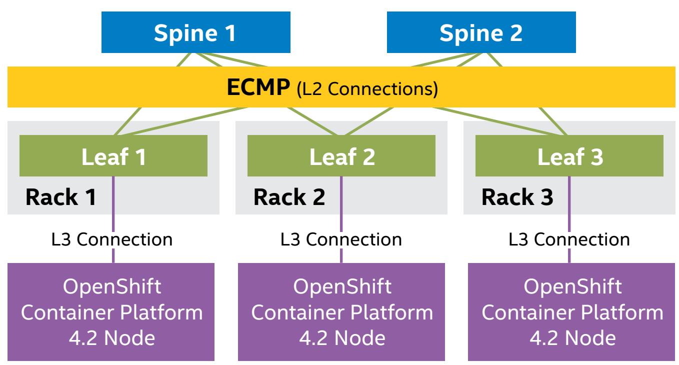

Figure 5. CDP Private Cloud network configuration recommendations.

Figure 6, on the next page, illustrates the networks that are used for the CDP Private Cloud Base and CDP Private Cloud Experiences clusters, including the interconnect. Connectivity between the clusters and existing network infrastructure can be adapted to specific installations.

A common scenario is when the cluster data network is exposed to an existing network. In this scenario, the edge network is either not used or is used for application access or ingest processing.

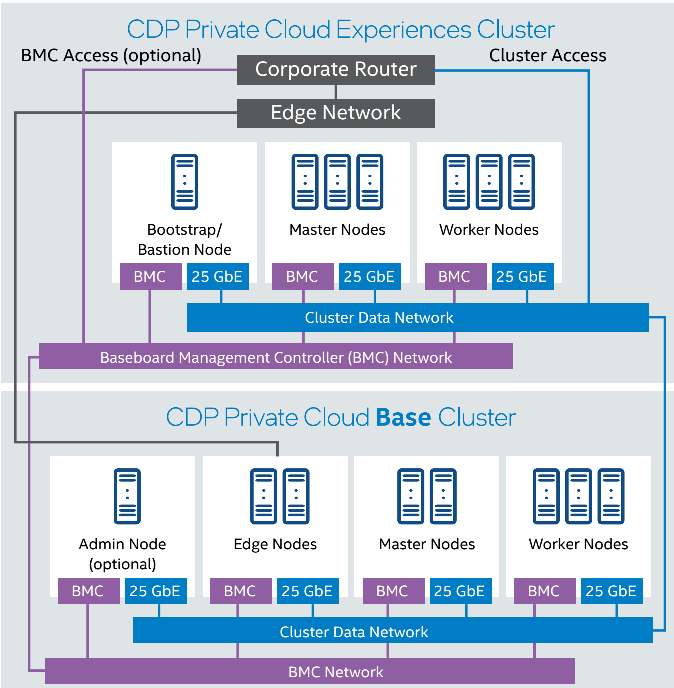

Figure 6. CDP Private Cloud cluster networks.

Network Functions

The CDP Private Cloud Cluster network functions include:

- **CDP Private Cloud Base Cluster Data Network.** The data network carries the bulk of the traffic within the cluster. This network is aggregated within each pod, and pods are aggregated into the cluster switch. The CDP Private Cloud Base services are available on this network. **Note:** The CDP Private Cloud Base services do not support multihoming and are only accessible on the cluster's data network.
- **CDP Private Cloud Experiences Cluster Data Network.** The data network carries the bulk of the traffic within the cluster. This network is aggregated within each POD (point of delivery), and PODs are aggregated into the cluster switch. The CDP Private Cloud Experiences services are available on this network.
- **Baseboard Management Controller (BMC) Network.** The BMC network connects the BMC ports and the out-of-band management ports of the switches. It is used for hardware provisioning and management. This network is aggregated into a management switch in each rack. This network provides access to the BMC functionality on the servers. It also provides access to the management ports of the cluster switches.
- **Edge Network.** The edge network provides connectivity from one or more edge nodes to an existing on-premises network, either directly or by the pod or cluster aggregation switches. SSH access to one or more edge nodes is available on this network, and other application services may be configured and available.

Cluster Interconnect Sizing

For cluster interconnect sizing information, see Cloudera's networking documentation.

Reference Solution Overview

Get better performance by upgrading to a newer version of Cloudera Data Platform with 2nd Gen Intel® Xeon® Scalable processors. Overall, this reference solution (see Figure 7) is an effective big data extension to an enterprise data warehouse analytics platform that offers the following:

- Scalability, agility, and flexibility
- Excellent performance and total cost of ownership
- Ability to satisfy business requirements for service-level agreements (SLAs), multiple users, and future growth

The solution can be applied to both green-field and refresh deployment scenarios and is designed to be used to its fullest extent before scaling out by adding more nodes.

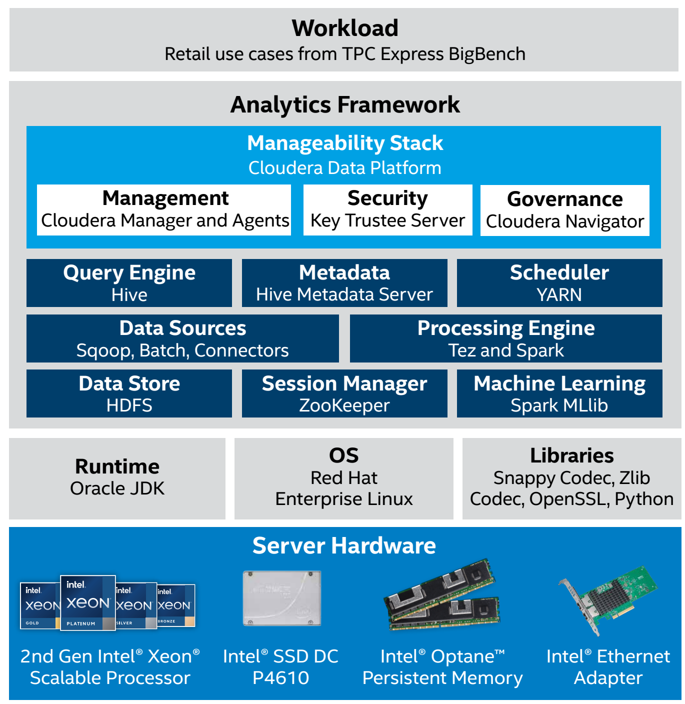

Figure 7. Reference solution for Cloudera Data Platform, running on Intel® hardware.

We have defined a highly available reference solution that includes the following nodes:

- **Edge nodes.** One edge node is for Cloudera Manager; the other two are master nodes (NameNode and Resource Manager). Edge nodes support services that are needed for the cluster operation. If one edge node fails, the other two can pick up the extra work.
- **Worker nodes.** This reference solution uses an Intel® SSD on every worker node. Worker nodes handle the bulk of the Hadoop processing. The number of worker nodes necessary depends on dataset size. Depending on your scalability needs, you can replicate this configuration to 20 or 30 worker nodes, or half it to four or five worker nodes. Some testing was performed with four worker nodes, but 10 worker nodes represent a nominal workload configuration.

Benchmark Testing: Intel® Technology Significantly Improves Big Data Analysis Performance

We ran two tests using Intel technology with CDP Private Cloud Base (each of these tests are described in detail in the following subsections):

- **BigBench Benchmark**. We compared Spark SQL performance on 2nd Generation Intel Xeon Scalable processors and CDP Private Cloud Base against the performance of a legacy Cloudera Distribution of Hadoop (CDH) on a previous-generation Intel Xeon processor.
- **Yahoo! Cloud Serving Benchmark** (YCSB). We compared Apache Kudu performance with and without Intel Optane PMem.

BigBench Results: Newer Software and Hardware Doubles Throughput

The BigBench benchmark is derived from the TPCx-BB Express Benchmark BB, an industry-standard benchmark licensed under TPC that measures the performance of Hadoop-based big data systems. BigBench measures the performance of both hardware and software components by executing 30 frequently performed analytical queries in the context of retailers with physical and online store presence (see Appendix B). The queries are expressed in SQL for structured data and in machine-learning algorithms for semi-structured and unstructured data. The SQL queries can use Hive or Spark, while the machine-learning algorithms use machine-learning libraries, user-defined functions, and procedural programs.

When we compared the performance (see Figure 8) of CDP Private Cloud Base version 7.0.3 running on the Intel Xeon Gold 6248 processor to CDH 5.16.1 running on the Intel Xeon processor E5-2680 v3, normalized big data batch analytics performance for the BigBench retail use cases improved by up to 2.23x for two streams and by up to 2.21x for four streams.[12] Table 8 provides the bill of materials for the tested cluster; see the appendices for additional configuration details.

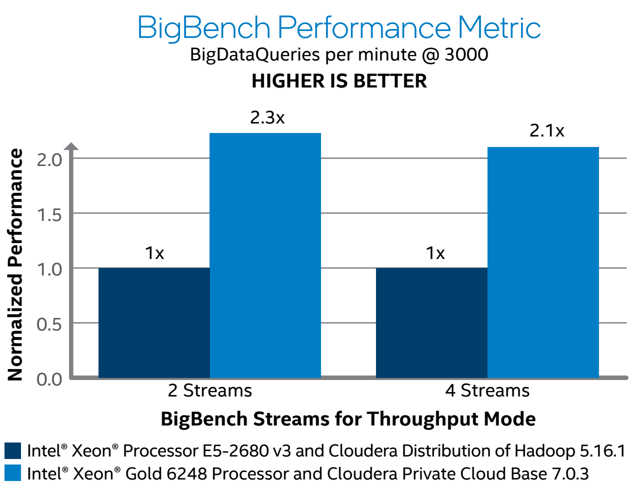

Figure 8. Double your throughput by upgrading to Cloudera Data Platform Private Cloud Base on the latest Intel® Xeon® Scalable processor and Intel® 3D NAND SSD.

Table 8. Bill of Materials for Tested Cluster

Description		Comments
PROCESSOR (per node)		
Management		Intel® Xeon® Gold 6248 processor @ 2.5/3.9 GHz • 1x Master Node/Active NameNode • 1x Utility Node/Standby NameNode • 1x Utility Node/Management Node
Worker		Intel Xeon Gold 6248 processor @ 2.5/3.9 GHz (at least 4 worker nodes; can scale as necessary)
Memory		12x 32 GB DDR4 @ 2666 MHz (total 384 GB per node)
Network		1x Intel® Ethernet Adapter X722 (10 GbE)
STORAGE		
Management Nodes (per node)		
	OS	Intel® SSD DC S4500 960 GB (total 8 TB per worker node)
	HDFS	2x 4 TB HDD 3.5-inch SAS3 12 Gb/s 7200 RPM Seagate Enterprise V.5 ST4000NM0095 (total 8 TB per management node)
Worker Nodes (per node)		
	OS	Intel® SSD DC S4500 960 GB (total 8 TB per worker node)
	HDFS	8x 4 TB HDD 3.5-inch SAS3 12 Gb/s 7200 RPM Seagate Enterprise V.5 ST4000NM0095 (total 32 TB per worker node)
	Fast Storage	Intel® SSD DC P4610 Series 1.6 TB, 2.5-inch PCIe 3.1 x4, 3D2, TLC (for YARN tmp files and Spark shuffles[a])

[a] The testing used Intel® SSD DC P4610 drives due to availability. We now recommend newer Intel® SSD DC P5510 drives (from the same product family).

The appropriate number of worker nodes in the cluster depends on the dataset size. For smaller datasets, four worker nodes may be sufficient; for larger datasets, you may need 10 worker nodes. For distributed processing, you could scale out to 20 or 30 worker nodes, if necessary. It is expected that an HDD plus one 3D NAND SSD per worker node is more cost effective compared to using a dense, all-SSD-based solution.

Apache Kudu: Intel Optane PMem Improves Throughput and Latency

Intel Optane PMem represents a new class of memory and storage technology. It is designed to improve the overall performance of the server by providing large amounts of persistent storage with low-latency access. Intel Optane PMem modules are DDR4-socket compatible and are offered in sizes not available with typical DDR4 DRAM products. The default reference design does not use PMem. However, the configuration can be upgraded to use PMem modules without any other additional hardware changes to gain a significant performance boost.

Intel Optane PMem can work in two different operating modes (Memory Mode and App Direct Mode), depending on business need and/or application support. In App Direct Mode, the DRAM and the PMem are both counted in the total system memory. This mode requires additional support from the OS and applications because there are two types of memory that can be used by the system independently. Low-latency operations should be directed to DRAM, while structures that are very large can be routed to PMem.

Intel Optane PMem has higher bandwidth and lower latency than SSD and HDD storage drives. These characteristics of PMem provide a significant performance boost to big data storage platforms that can utilize it for caching. One of such platforms is Apache Kudu—an open source columnar storage engine that enables fast analytics on fast data. Kudu Tablet Servers store and deliver data to clients.

Kudu can utilize PMem for its internal block cache. We tested Kudu performance using YCSB. We compared how Kudu performs with block cache in DRAM to learn how it performs when using PMem for block cache. To obtain maximum performance for Kudu block cache implementation, we used the Persistent Memory Development Kit (PMDK) and PMem in App Direct Mode.

Using PMem for block cache confers several benefits:

- Reduce the DRAM footprint required for Kudu
- Keep performance as close to DRAM speed as possible
- Take advantage of larger cache capacity to cache more data and improve the entire system's performance

We tested two datasets: small (100 GB) and large (700 GB). The small dataset is designed to fit entirely inside the Kudu block cache for both DRAM and PMem. The large dataset is designed to exceed the capacity of Kudu block cache on DRAM, while fitting entirely inside the block cache on PMem. We ran same the YCSB workload for each dataset, with five iterations. The first iteration is a warm-up phase to load data from disk to cache; the other iterations read data from the cache. See Appendix B for Kudu and YCSB parameter configurations.

For the large dataset, the PMem-based configuration delivered up to a 6.3x increase in throughput compared to the all-DRAM-based configuration. When measuring latency of reads at the 95th percentile (reads with observed latency higher than 95 percent of all other latencies), we observed a 13.38x decrease in latency on the PMem-based configuration compared to the all-DRAM-based configuration (see Figure 9). For the small dataset we observed similar performance in both configurations—that is, PMem performance was nearly identical to all-DRAM performance.[13]

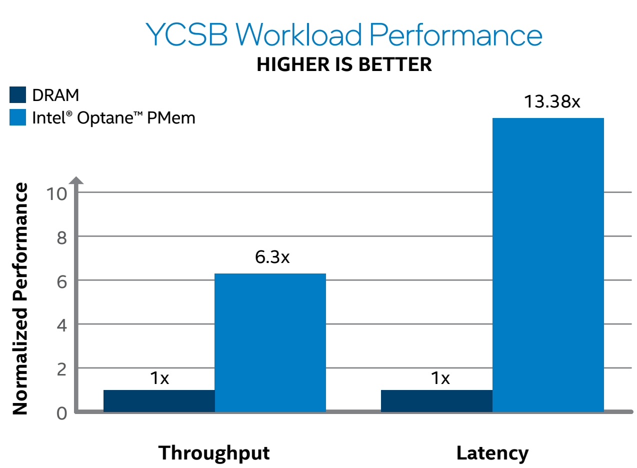

Figure 9. Using Intel® Optane™ PMem for Kudu block cache for large datasets that exceed DRAM capacity can increase throughput by more than 6x and cut latency by over 13x.

Migration Path from Legacy Cloudera Distributions to CDP Private Cloud

Cloudera no longer supports Hortonworks Data Platform (HDP) 2.6.x and CDH 5.x. Therefore, customers using these products need to map a path for upgrading to a supported version.

Customers using Cloudera Enterprise 5.13+ and HDP 2.6.5, or Cloudera Enterprise 6 and HDP 3, should first migrate their workloads to CDP Private Cloud Base 7.1.3+ (a shared data experience environment). Existing applications, data, and hardware can be brought into the new environment; there is also opportunity for hardware refreshes that can improve performance as noted in the Solution Architecture section. Once the migration to CDP Private Cloud Base is complete, then organizations can choose to expand to new experiences by upgrading to CDP Private Cloud Experiences (see Figure 10).

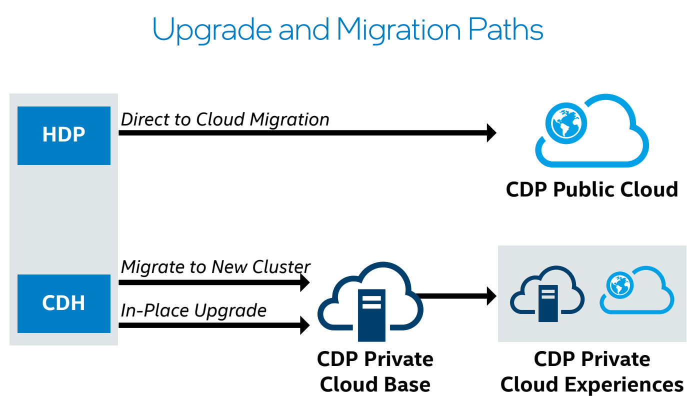

Figure 10. Hortonworks Data Platform or Cloudera Enterprise Data Hub users can follow one of the several upgrade or migration paths to CDP.

Before upgrading a cluster to CDP Private Cloud Base, the existing infrastructure should be evaluated to determine whether it meets the needs of a full CDP Private Cloud Base environment. For in-place upgrades, the primary requirement is that the cluster infrastructure should meet the same requirements as those specified for CDH or HDP in the Cloudera Reference Architecture Documentation.

Beyond those recommendations, our general guidance is:

- After CDP Private Cloud is deployed, most of the CDP Private Cloud Base compute workloads will begin migrating to the new infrastructure. The CDP Private Cloud Base infrastructure becomes a storage-centric cluster. This configuration means that smaller memory and lighter processor configurations may be adequate for the CDP Private Cloud Base cluster.
- CDP Private Cloud relies heavily on the network infrastructure for data access. We recommend a minimum of two bonded 10 GbE connections, with 25 GbE or faster preferred. Clusters with 1 GbE networks should probably not be upgraded. Although faster network cards can be installed, consider the infrastructure and labor costs.
- If the cluster is due for an infrastructure refresh, we suggest deploying a new CDP Private Cloud environment and then migrating data and workloads to the new environment.

Additional Infrastructure Considerations

This section provides information about infrastructure topics that may be useful as organizations deploy CDP Private Cloud.

Repurposing Infrastructure

Repurposing infrastructure is a common topic in upgrade conversations; this document provides some suggestions based on experience.

The recommended CDP Private Cloud Base worker nodes are storage-heavy configurations, while the CDP Private Cloud Experiences worker nodes are memory- and compute-heavy. Older CDH, HDP, or CDP-DC worker nodes generally cannot be reused in the CDP Private Cloud Experiences cluster for the cloud workloads. Those nodes are potentially useful in another environment for storage-heavy applications. In some instances, it may be possible to add them to the OpenShift cluster for non-CDP Private Cloud Experiences workloads.

After a cluster refresh, we often see that customers repurpose older CDH or HDP nodes for development, test, or disaster recovery clusters. In many instances, those nodes are added to existing clusters that are not performance-critical.

Heterogenous Nodes

The core OpenShift platform supports heterogenous node types, including compute-, memory-, and accelerator-optimized configurations. The necessary support for seamless use of heterogenous nodes is not currently available in either OpenShift 4.2 or CDP Private Cloud Experiences. They are not recommended for initial deployments. Software support in this area is evolving rapidly, so if these configurations are of interest, contact your Intel or Cloudera representative for the latest status.

Using an Existing OpenShift Cluster

Depending on the OpenShift node configurations, this reuse may be possible. Contact your Intel or Cloudera representative for the latest status.

Ozone Object Storage

Ozone is a scalable, redundant, and distributed object store optimized for big data workloads. Ozone can scale to billions of objects of varying sizes and can function effectively in containerized environments such as Kubernetes and YARN.

We recommend using at least one Intel® Optane™ SSD DC P5800X Series (800 GB, 2.5in PCIe x4) per node, and expect that adding another of these next-generation SSDs for the Ozone cache will further improve performance. From an infrastructure planning perspective, we expect the initial release of Ozone to run on the CDP Private Cloud Base cluster and use storage resources from that infrastructure. For customers interested in Ozone, visit Cloudera's Ozone web page.

Sizing CDP Private Cloud Base Clusters

Cluster nodes are broadly described as master nodes, utility nodes, gateway nodes, or worker nodes.

- Master nodes run Hadoop master processes such as the HDFS NameNode and YARN Resource Manager.
- Utility nodes run other cluster processes that are not master processes, such as Cloudera Manager and the HMS.
- Gateway nodes are client access points for launching jobs in the cluster. The number of gateway nodes required varies depending on the type and size of the workloads.
- Worker nodes primarily run DataNodes and other distributed processes such as Impala.

Cloudera provides guidance on sizing clusters, recommending role allocations for different cluster sizes ranging from three to 10 worker nodes, up to 500 to 1,000 worker nodes.

Note: Cloudera recommends that you always enable high availability when using Runtime in a production environment.

Learn More

You may find the following references helpful:

Intel

- 2nd Generation Intel® Xeon® Scalable processors
- Intel® Optane™ persistent memory
- Intel® Optane™ solid state drives
- Intel® Solid State Drives Data Center Family
- Intel® Ethernet Network Adapter E810-XXVDA2 for OCP 3.0 (dual-port 25 GbE)

Cloudera

- Cloudera Data Platform (CDP) Private Cloud

Find the solution that is right for your organization. Visit **intel.com/AI** or contact an Intel representative.

Appendix A: Additional Configuration Details for BigBench Testing

Tables A1-A3 provide configuration details for the BigBench cluster testing. Table A4 shows which CDP Private Cloud roles are present on which nodes. For more information, read the Cloudera article on this topic.

Table A1. Software Configuration

	Intel® Xeon® Processor E5-2680 v3	Intel® Xeon® Gold 6248 Processor
OS	RHEL 7.6	RHEL 7.8
Kernel	3.10.0-957.12.2.el7.x86_64	3.10.0-1127.10.1.el7.x86_64
Java	jdk1.7.0_67-cloudera	jdk1.8.0_232-cloudera
Workload	BigBench (based on TPCx-BB v1.3)	BigBench (based on TPCx-BB v1.3.1)
Cloudera Distribution	CDH 5.16.1	Cloudera Data Platform Private Cloud Base 7.0.3

Table A2. Benchmark Configuration

Benchmark Configuration	
Total Nodes	• 5x (1x Name-node, 4 Data-nodes) for Intel® Xeon® processor E5-2680 v3 – 3 TB Scale Factor • 7x (1x Master + 2x Management [recommended by Cloudera] + 4x Worker Nodes) for Intel® Xeon® Gold 6248 processor – 3 TB Scale Factor • 13x (1x Master + 2x Management [recommended by Cloudera] + 10x Worker Nodes) for Intel Xeon Gold 6248 processor – 10 TB Scale Factor
Workload	BigBench
Storage	HDFS, RF = 3
SQL Engine	Hive on Spark for CDH 5.16, SparkSQL for CDP Base 7.0.3 (Hive on Spark no longer supported by Cloudera)
Scale Factor	3 TB and 10 TB for BigBench
Streams	2 and 4 Parallel Streams

Table A3. Spark Configuration

		Intel® Xeon® Processor E5-2680 v3	Intel® Xeon® Gold 6248 Processor
Number of Executors		9	16
Executor Size (Heap Size)	**Cores**	5	5
	Memory	25 GB	22.5 GB

Table A4. Cloudera Data Platform Private Cloud Role Distribution

Cloudera Role	Master Node 1	Master Node 2	Edge Node	Worker Nodes
HDFS NameNode	Y			
HDFS Secondary NameNode		Y		
HDFS DataNode				Y
YARN Resource Manager	Y			
YARN Job History Server	Y			
YARN Node Manager				Y
HiveServer2	Y			
Hive Metastore Server		Y		
ZooKeeper	Y	Y	Y	
Hive Gateway	Y	Y	Y	Y
Spark Gateway	Y	Y	Y	Y
Spark History Server	Y			
Cloudera Manager			Y	
Cloudera Manager Management Service			Y	
Apache Ranger	Y			
Apache Atlas	Y			

Appendix B: BigBench Use Case Descriptions

Refer to the following resources for descriptions of the parameters listed in Table B1:

- https://hadoop.apache.org/docs/r3.0.0/hadoop-mapreduce-client/hadoop-mapreduce-client-core/mapred-default.xml
- https://cwiki.apache.org/confluence/display/Hive/Configuration+Properties

Table B1. Workload Use Cases and Optimizations

	Use Case	Primary Data Type
1	Find top 100 products that are sold together frequently in given stores.	Structured
2	Find the top 30 products that are mostly viewed together with a given product in online store.	Semi-structured
3	For a given product, get a top-30 list sorted by number of views in descending order of the last five products that are mostly viewed before the product was purchased online.	Semi-structured
4	Shopping cart abandonment analysis: For users who added products in their shopping carts but did not check out in the online store during their session, find the average number of pages they visited during their sessions.	Semi-structured
5	Build a model using logistic regression for a visitor to an online store, based on existing users' online activities (interest in items of different categories) and demographics.	Semi-structured
6	Identify customers shifting their purchase habit from physical store to Web sales.	Structured
7	List top-10 states in descending order with at least 10 customers, who during a given month bought products with the price at least 20 percent higher than the average price of products in the same category.	Structured
8	For online sales, compare the total sales amount for which customers checked online reviews before making the purchase and that of sales for which customers did not read reviews. Consider only online sales for a specific category in a given year.	Semi-structured
9	Aggregate total amount of sold items over different combinations of customers based on selected groups of marital status, education status, sales price, and different combinations of state and sales profit.	Structured
10	For all products, extract sentences from reviews that have positive or negative sentiment.	Unstructured
11	For a given product, measure the correlation of sentiments, including the number of reviews and average review ratings, on product monthly revenues within a given time frame.	Semi-structured
12	Find all customers who viewed items of a given category on the Web in a given month and year that was followed by an in-store purchase of an item from the same category in the three consecutive months.	Semi-structured
13	Display customers with both store and Web sales in consecutive years for which the increase in Web sales exceeds the increase in-store sales for a specified year.	Structured
14	Calculate the ratio between the number of items sold on the Web in the morning (7 to 8 AM) to the number of items sold in the evening (7 to 8 PM) for customers with a specified number of dependents.	Structured
15	Find the categories with flat or declining sales for in-store purchases during a given year for a given store.	Structured
16	Compute the impact of an item price change on store sales by computing the total sales for items in a 30-day period before and after the price change.	Structured
17	Find the ratio of items sold with and without promotions in a given month and year. Only items in certain categories sold to customers living in a specific time zone are considered.	Structured
18	Identify the stores with flat or declining sales in three consecutive months, and check if there are any negative reviews regarding these stores available online. Analyze the online reviews for these items to determine if there are any major negative reviews.	Unstructured
19	Retrieve the items with the highest number of returns, where the number of returns was approximately equivalent across all store and Web channels.	Unstructured
20	Perform customer segmentation for return analysis. Segment customers according to varying criteria such as return frequency, returns to order ratio, and return amount ratio.	Structured
21	Get all items that were sold in stores in a given month and year, and which were returned in the next six months and repurchased by the returning customer afterwards through the Web sales channel in the following three years.	Structured
22	Compute the percentage change in inventory between the 30-day period before the price change and the 30-day period after the change.	Structured
23	Calculate the coefficient of variation and mean of every item and warehouse of the given and the consecutive months.	Structured
24	For a given product, measure the effect of competitors' prices on the product's in-store and online sales.	Structured
25	Perform customer segmentation analysis. Segment customers according to key shopping dimensions such as how recent the last visit was, frequency of visits, and monetary amount.	Structured
26	Cluster customers into book buddies or groups based on their in-store book purchasing histories.	Structured
27	Extract competitor product names and model names (if any) from online product reviews for a given product.	Unstructured
28	Build text classifier for online review sentiment classification (Positive, Negative, Neutral).	Unstructured
29	Perform category affinity analysis for products purchased together online.	Structured
30	Perform category affinity analysis for products viewed together online.	Semi-structured

NLP – natural language processing; **UDF** – user-defined function; **UDTF** – user-defined table function

Appendix C: Benchmark Tuning Parameters

This appendix provides the detailed tuning parameters used in the benchmark tests described earlier in this document.

BigBench Tuning Parameters

Cloudera Data Platform Private Cloud Base includes many default parameter settings. Table C1 shows only the tuning parameters we changed from the default setting to achieve the best performance from the testing cluster.

Table C1. BigBench Tuning Parameters

Property	Value
YARN	
yarn.scheduler.maximum-allocation-mb	360 GB
yarn.scheduler.maximum-allocation-vcores	80
yarn.nodemanager.resource.memory-mb	360 GB
yarn.nodemanager.resource.cpu-vcores	80
Client Java Heap Size in Bytes	4 GB
Hadoop Distributed File System (HDFS)	
dfs.namenode.handler.count	200
dfs.namenode.service.handler.count	200
dfs.datanode.handler.count	30
dfs.permissions	FALSE
dfs.datanode.socket.write.timeout	630000
dfs.socket.timeout	630000
dfs.datanode.max.transfer.threads	65536
Dfs.blocksize	256 MB
MapReduce	
mapreduce.output.fileoutputformat.compress	TRUE
mapreduce.output.fileoutputformat.compress.codec	org.apache.hadoop.io.compress.SnappyCodec
mapreduce.map.output.compress.codec	org.apache.hadoop.io.compress.SnappyCodec
mapreduce.map.memory.mb	4.5 GB
mapreduce.reduce.memory.mb	4.5 GB
mapreduce.map.java.opts.max.heap	4 GB
mapreduce.reduce.java.opts.max.heap	4 GB
mapreduce.reduce.shuffle.parallelcopies	80
mapreduce.job.reduce.slowstart.completedmaps	0.9
Io.sort.mb	512
Spark	
spark.master	yarn-client
spark.network.timeout	9000 seconds
spark.eventLog.enabled	TRUE
spark.eventLog.dir	hdfs:///user/spark/applicationHistory
spark.serializer	org.apache.spark.serializer.KryoSerializer
spark.default.parallelism	630 (3 TB) / 2385 (10 TB)
spark.kryo.referenceTracking	FALSE
spark.io.compression.codec	lzf
spark.driver.extraJavaOptions	-XX:+UseParallelOldGC -XX:ParallelGCThreads=5 -XX:NewRatio=1 -XX:SurvivorRatio=1 -XX:+UseCompressedOops
spark.executor.extraJavaOptions	-XX:+UseParallelOldGC -XX:ParallelGCThreads=5 -XX:NewRatio=1 -XX:SurvivorRatio=1 -XX:+UseCompressedOops
spark.yarn.maxAppAttempts	1
spark.driver.memory	200 GB
spark.executor.cores	5
spark.executor.memory	18475 MB
spark.yarn.executor.memoryOverhead	4096 MB
Spark SQL	
spark.sql.shuffle.partitions	630 (3 TB) / 2385 (10 TB)
spark.sql.statistics.size.autoUpdate.enabled	TRUE
spark.sql.autoBroadcastJoinThreshold	536870912
spark.sql.broadcastTimeout	5400
spark.sql.cbo.enabled	TRUE
spark.sql.warehouse.dir	/warehouse/tablespace/managed/hive

Property	Value
SparkSQL (only for 10 TB)	
spark.sql.hive.caseSensitiveInferenceMode	NEVER_INFER
spark.sql.hive.metastorePartitionPruning	FALSE
spark.sql.hive.convertMetastoreOrc	FALSE

Apache Kudu and Yahoo! Cloud Serving Benchmark (YCSB) Tuning Parameters

Cloudera Data Platform Private Cloud includes many default parameter settings. Table C2 provides the YCSB workload properties we used in testing, while Table C3 provides our Kudu configuration settings. Table C4 provides details on the YCSB dataset.

Table C2. YCSB Workload Properties

Property	Value
readallfields	TRUE
readproportion	1
updateproportion	0
scanproportion	0
insertproportion	0
Fieldlength (700 GB dataset)	30000
Fieldcount (700 GB dataset)	200
Fieldlength (100 GB dataset)	10000
Fieldcount (100 GB dataset)	100
Requestdistribution	sequential

Table C3. Apache Kudu Configuration

Parameter	All-DRAM Cluster	DRAM+PMem Cluster
unlock_unsafe_flags	TRUE	TRUE
force_block_cache_capacity	TRUE	TRUE
maintenance_manager_num_threads	4	4
memory_limit_hard_bytes	768000000000	192000000000
rpc_service_queue_length	100	100
rpc_authentication	disable	disable
rpc-encryption	disable	disable
unlock_experimental_flags	TRUE	TRUE
block_cache_type	DRAM	NVM
block_cache_capacity_mb	720000	920000

Table C4. YCSB Dataset Properties

Data size (on disk)	700 GB	100 GB
In-memory data	~900 GB	~110 GB

Solution Provided By:

1 Tested by Intel June 20-24, 2019 (baseline configuration) and May 27-June 27, 2020 (upgraded configuration).

Baseline configuration: 4 worker nodes and 1 management node.

System under test: 2x Intel® Xeon® processor E5-2680 v3 (12 cores, 24 threads), microcode = 0x43; Intel® Hyper-Threading (Intel® HT) Technology = ON; Intel® Turbo Boost Technology (Intel® TBT) = ON; BIOS = SE5C610.86B.01.01.0028.121720182203; DRAM = 24 slots/16 GB/2133 MHz (total memory = 384 GB); boot drive = 1x Intel® DC SSD S3500 120 GB (3 nodes) and 1x Intel DC SSD S3500 240 GB (1 node); application storage = 8x 3 TB 7200K RPM HDD; network = 1x Intel® Ethernet Adapter X540-AT2 (10 GbE).

Management node: Intel® Xeon® processor E5-2699 v3 (18 cores, 36 threads), microcode = 0x43; Intel HT = ON; Intel TBT = ON; DRAM = 24 slots/16 GB/2133 MHz (total memory = 384 GB); storage = 2x 3 TB 7200K RPM HDD; network = 1x Intel Ethernet Adapter X540-AT2 (10 GbE).

Upgraded configuration: 4 or 10 worker nodes and 3 management nodes.

System under test: 2x Intel® Xeon® Gold 6248 processor (20 cores, 40 threads), microcode = 0x500002c; Intel HT = ON; Intel TBT = ON; BIOS = SE5C620.86B.02.01.0008.031920191559; DRAM = 12 slots/32 GB/2666 MHz (total memory = 384 GB); boot drive = 960 GB Intel® SSD D3-S4610 (SATA); application storage = 8x 4 TB 7200 RPM 12 Gb/s Seagate Enterprise Capacity 3.5'' HDD 4TB 7200 RPM and 1x 1.6 TB Intel® DC SSD P4610 (3D NAND QLC) for Spark shuffle; network = 1x Intel® Ethernet Adapter X722 (10 GbE).

Management nodes: 2x Intel Xeon Gold 6248 processor (20 cores, 40 threads), microcode = 0x500002c; Intel HT = ON; Intel TBT = ON; DRAM = 12 slots/32 GB/2666 MHz (total memory = 384 GB); storage = 2x 4 TB 7200 RPM 12 Gb/s Seagate Enterprise Capacity 3.5'' HDD 4TB 7200 RPM; network = 1x Intel Ethernet Adapter X722 (10 GbE).

2 Tested by Intel May 2021.

Common software and workload configuration: OS = CentOS Linux 7 (Core); kernel = 3.10.0-1160.25.1.el7.x86_64; workload and version = YCSB 0.17.0 – pure read.

Common test methodology: run method = warm, 6 iterations (1 warm-up, 5 read-from-cache); result choice = median; dataset size = 700 GB and 100 GB

All-DRAM configuration: Single node; Intel® Server Board S2600WFT; 2x Intel® Xeon® Gold 6248R processor (24 cores; 3.0 GHz); microcode = 0x5003006; Intel HT = ON; Intel TBT = ON; BIOS = SE5C620.86B.02.01.0008.031920191559; total memory = 768 GB (24 slots/32 GB/2666 MHz); storage (boot) = 1.92 TB Intel® SSD DC S4600; storage (application drives) = 2x 1.92 TB Intel SSD DC S4600.

DRAM + Intel® Optane™ PMem configuration: Single node; Intel® Server Board S2600WFT; 2x Intel Xeon Gold 6248R processor (24 cores; 3.0 GHz); microcode = 0x5003006; Intel HT = ON; Intel TBT = ON; BIOS = SE5C620.86B.02.01.0008.031920191559; DDR memory = 384 GB (12 slots/32 GB/2666 MHz); Intel® Optane™ persistent memory (PMem) = 1024 GB (8 slots/128 GB/QS); storage (boot) = 1.92 TB Intel SSD DC S4600; storage (application drives) = 2x 1.92 TB Intel SSD DC S4600; DCPM firmware version = 01.02.00.5417; PMem mode = App Direct.

3 Source: gartner.com/en/newsroom/press-releases/2018-02-05-gartner-survey-shows-organizations-are-slow-to-advance-in-data-and-analytics

4 See endnote 1.

5 See endnote 2.

6 cloudera.com/about/customers/iqvia.html

7 See endnote 1.

8 See endnote 2.

9 See endnote 1.

10 See endnote 2.

11 For more information about CDP Private Cloud, visit cloudera.com/products/cloudera-data-platform/cdp-private-cloud.html.

12 See endnote 1.

13 See endnote 2.

See Appendix A and Appendix C for additional configuration information.