



Cisco Data Intelligence Platform on Cisco UCS M6 with Cloudera Data Platform Ozone Design Guide

Published: February 2022



In partnership with:

CLOUDERA



About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_U1).

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2022 Cisco Systems, Inc. All rights reserved.

Contents

Executive Summary	6
Solution Overview	8
Technology Overview	19
Solution Design	50
Solution Prerequisites	55
Summary	59
Bill of Materials	61
About the Authors	64
Feedback	65

Executive Summary

Enterprise data analytics teams are always seeking for new ways to improve their systems. Storage is a critical component of the data platforms strategy since it serves as the foundation for all compute engines and applications that are developed on top of it. Businesses are also considering switching to a scale-out storage strategy that offers dense storage as well as reliability, scalability, and performance. To make this a reality, Cisco and Cloudera collaborated on testing with dense storage nodes.

Over time, the Hadoop ecosystem has evolved from batch processing (Hadoop 1.0) to streaming and near real-time analytics (Hadoop 2.0), and finally to Hadoop meets AI (Hadoop 3.0). The capabilities of the technologies are now evolving to enable data lake as a private cloud with the storage and compute isolation, as well as enabling Hybrid cloud in the future (and multi-cloud).

Cloudera released the following software in the second half of 2020, both of which enable the data lake as a private cloud:

- **Cloudera Data Platform Private Cloud Base (CDP PvC Base)** – provides storage and supports the traditional Data Lake environments and introduced Apache Ozone, the next generation filesystem for Data Lake
- **Cloudera Data Platform Private Cloud Data Services (CDP PvC DS)** – provide different experiences or personas (data analyst, data scientist, data engineer) driven data experiences from private and hybrid data lakes

Apache Ozone initiative provides the foundation for the next generation of storage architecture for Data Lake, where data blocks are organized in storage containers for higher scale and handling of small objects in HDFS. The Ozone project also includes an object store implementation to support several new use cases.

[Cisco Data Intelligence Platform](#) (CDIP) is thoughtfully designed private cloud for data lake. It supports data intensive workloads with Cloudera Data Platform Private Cloud Base, and compute rich (AI/ML) and compute intensive workloads with Cloudera Data Platform Private Cloud Data Services. CDIP further provides storage consolidation with Apache Ozone on Cisco UCS infrastructure, which is fully managed by Cisco Intersight. Cisco Intersight simplifies management and moves management of computing resources from the network into the cloud.

This CVD includes CDIP with cloud advantage in mind for private and hybrid cloud. It is based on Cisco UCS M6 family of servers which support 3rd Gen Intel Xeon Scalable family processors with PCIe Gen 4 capabilities. These servers include the following. These servers include the following.

- The Cisco UCS C240 M6 Server for Storage (Apache Ozone and HDFS) – Extends the capabilities of the Cisco UCS rack server portfolio with 3rd Gen Intel Xeon Scalable Processors supporting more than 43 percent more cores per socket and 33 percent more memory when compared with the previous generation.

Furthermore, with Cisco Intersight you get all the benefits of SaaS delivery and full life cycle management of network and compute. This empowers you to analyze, update, fix, and automate your environment in ways that were not possible before.

This CVD, includes [Cloudera Private Cloud Base](#) (CDP PvC Base) version 7.1.7 running Apache Ozone on Cisco Data Intelligence Platform consisting of Cisco UCS M6 family of servers which support 3rd Gen Intel Xeon Scalable family processors with PCIe Gen 4 capabilities.

CDIP with Cloudera Data Platform enables the customer to independently scale storage and computing resources as needed while offering an exabyte scale architecture with low total cost of ownership (TCO). It offers future-proof architecture with the latest technology provided by Cloudera.

Solution Overview

Introduction

Big Data and machine learning technologies have advanced to the point that they are now being used in production systems that run 24 hours a day, seven days a week. The intake, processing, storage, and analysis of data, as well as the seamless distribution of the output, findings, and insights of the analysis, all require a proven, trustworthy, and high-performance platform.

This solution implements Cloudera Data Platform Private Cloud Base (CDP PvC Base) big data storage solution optimized for Apache Ozone on Cisco Data Intelligence Platform (CDIP), a world-class platform specifically designed for demanding workloads that is both easy to scale and easy to manage, even as the requirements grow to thousands of servers and petabytes of storage.

Recognizing the enormous potential of big data and machine learning technologies, many enterprises are preparing to take advantage of these new capabilities by expanding departments and recruiting more people. However, these initiatives are up against a new set of obstacles:

- Making the data available to the diverse set of engineers (Data engineers, analysts, data scientists) who need it
- Enabling access to high-performance computing resources, GPUs, that also scale with the data growth
- Allowing people to work with the data using the environments in which they are familiar
- Publishing their results so the organization can make use of it
- Enabling the automated production of those results
- Managing the data for compliance and governance
- Scaling the system as the data grows
- Managing and administering the system in an efficient, cost-effective way

This solution is built on the Cisco Data Intelligence Platform, which includes compute, storage, connectivity, and unified management capabilities built on Cisco Unified Computing System (Cisco UCS) infrastructure, using Cisco UCS C-Series, X-Series and S-Series Rack Servers. Unified management with Cisco Intersight to help companies manage the entire infrastructure from a single pane of glass, as well as Cloudera Data Platform to provide the software for fast ingest of data and managing and processing it. This architecture is optimized for big data and machine learning workloads in terms of performance and linear scalability.

Audience

The intended audience of this document includes sales engineers, field consultants, professional services, IT managers, partner engineering and customers who want to deploy the Cloudera Data Platform on the Cisco Data Intelligence Platform on Cisco UCS M6 Rack-Mount servers and Cisco UCS® X-Series.

Purpose of this Document

This document describes the architecture, design choices, and deployment procedures for Cisco Data Intelligence Platform using Cloudera Data Platform Private Cloud Base on Cisco UCS M6 Rack-Mount servers.

What's New in this Release?

This solution extends the portfolio of Cisco Data Intelligence Platform (CDIP) architecture by combining Cloudera Data Platform Private Cloud with Apache Ozone, a cutting-edge distributed storage system that succeeds Hadoop Filesystem (HDFS). Apache Ozone solves some of HDFS' storage restrictions by offering native support for Objects and the S3 API, as well as greater storage consolidation per node and a smaller datacenter footprint. Furthermore, as the enterprise's demands and requirements evolve over time, the platform may scale to thousands of servers, exabytes of storage, and tens of thousands of cores to handle the data.

The following will be implemented in this validated design:

- Cisco Intersight to configure and manage Cisco Infrastructure
- Data lake provided by Cloudera Data Platform Private Cloud Base on Cisco UCS Servers
- Next-generation data storage with Apache Ozone providing both HDFS and S3 API support for data lake
- Intel based 3.8 TB U.2 P5500 series NVMe as metadata storage for Apache Ozone

Solution Summary

This CVD details the process of installing Cloudera Data Platform Private Cloud Base including the installation of Apache Ozone, the prerequisites for Cloudera Data Platform Private Cloud base, and the configuration details of the cluster.

Cisco Data Intelligence Platform

Cisco Data Intelligence Platform (CDIP) is a cloud-scale architecture and a private cloud primarily for a data lake which brings together big data, AI/compute farm, and storage tiers to work together as a single entity while also being able to scale independently to address the IT issues in the modern data center. This architecture provides the following:

- Extremely fast data ingest, and data engineering done at the data lake.

- AI compute farm allowing for different types of AI frameworks and compute types (GPU, CPU, FPGA) to work on this data for further analytics.



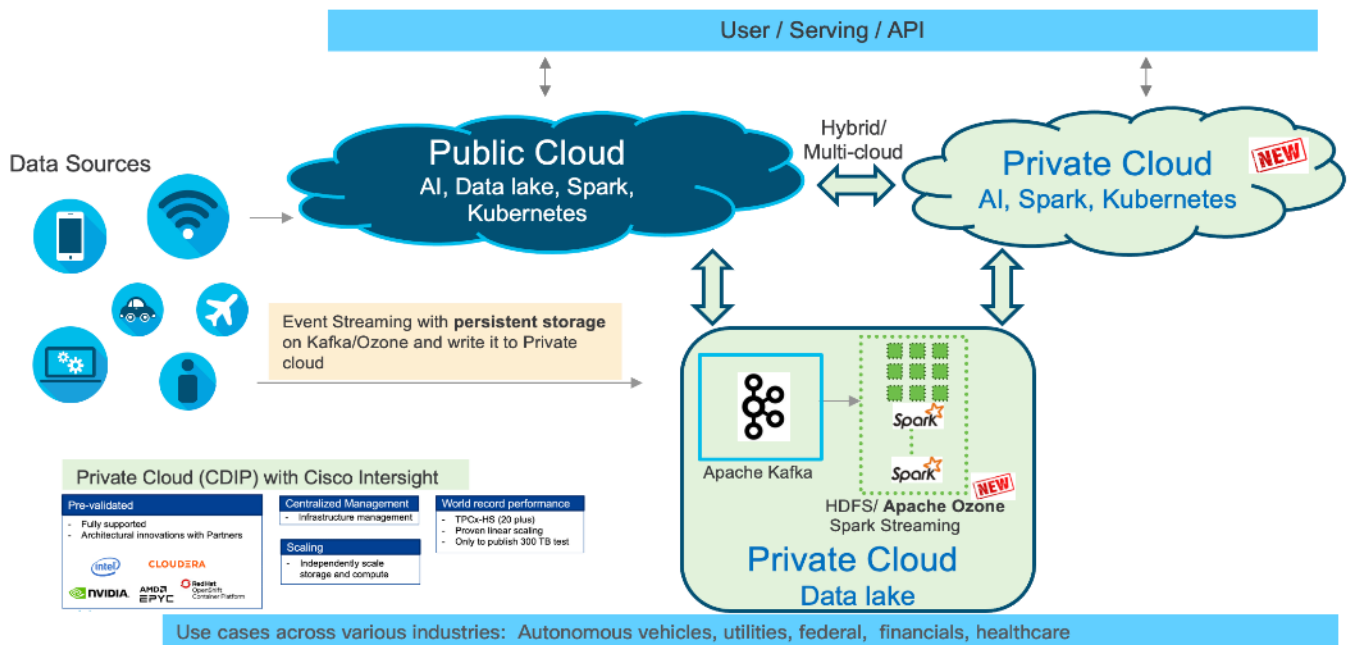
GPU and FPGA are not supported in this release of Cloudera Private Cloud Data Services 1.3

- A storage tier, allowing to gradually retire data which has been worked on to a storage dense system with a lower \$/TB providing a better TCO. Next-generation Apache Ozone filesystem for storage in a data lake.
- Seamlessly scale the architecture to thousands of nodes with a single pane of glass management using Cisco Intersight and Cisco Application Centric Infrastructure (ACI).

Cisco Data Intelligence Platform caters to the evolving architecture bringing together a fully scalable infrastructure with centralized management and fully supported software stack (in partnership with industry leaders in the space) to each of these three independently scalable components of the architecture including data lake, AI/ML and Object stores.

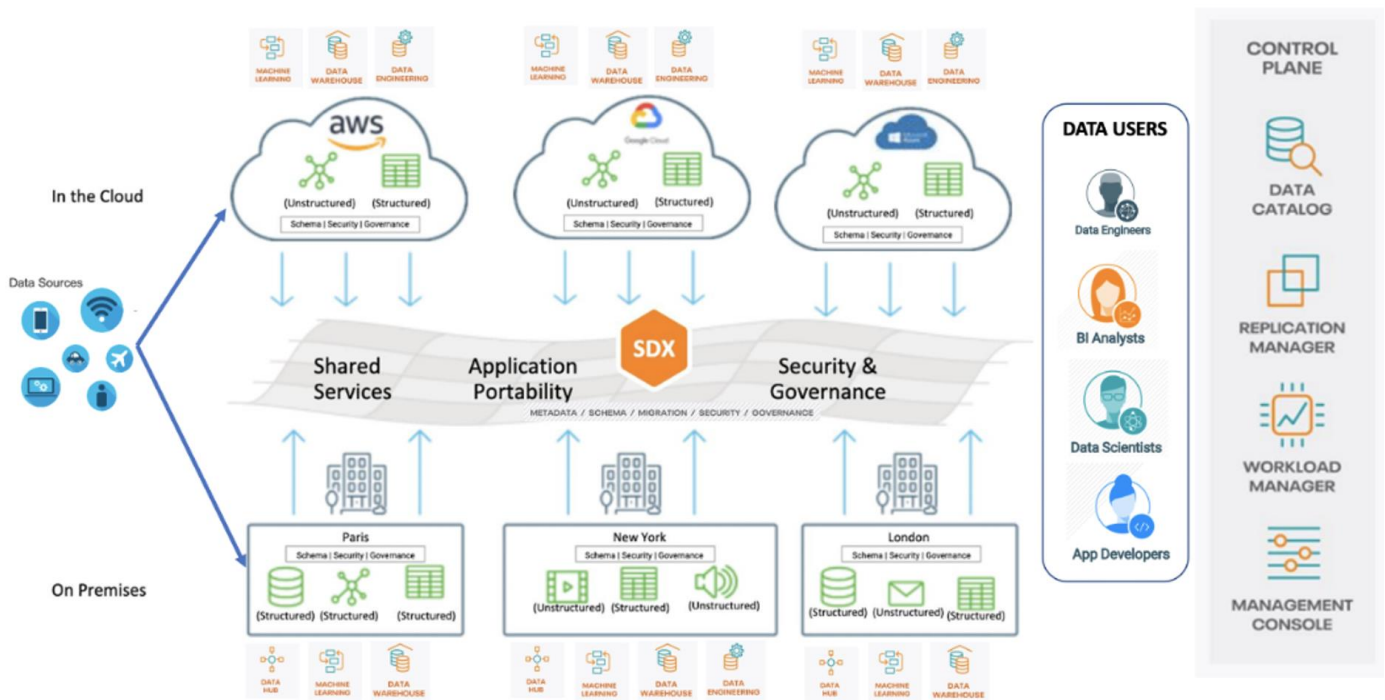
Figure 1. Cisco Data Intelligent Platform with Cloudera Data Platform - Evolution towards Hybrid Cloud

Cisco Data Intelligence Platform(CDIP) Evolution towards Hybrid Cloud



CDIP offers private cloud which enables it to become a hybrid cloud for their data lakes and apps which provides unified user experiences with common identity, single API framework that stretches from private cloud to public cloud, auto-scales when app demand grows. Further, implement tighter control over sensitive data with data governance and compliance, and integrate common data serving layer for data analytics, business intelligence, AI inferencing, and so on.

Figure 2. Cisco Data Intelligent Platform with Cloudera Data Platform - Hybrid Cloud Architecture



CDIP with CDP private cloud is built to meet the needs of enterprises for their hybrid cloud with unmatched choices such as any data, any analytics, and engineering anywhere. This solution includes:

- **Flexibility** to run workload anywhere for quick and easy insights.
- **Security** that is consistent across all clouds provided by Cloudera’s SDX. Write centrally controlled compliance and governance policies once and apply everywhere, enabling safe, secure, and compliant end-user access to data and analytics.
- **Performance and scale** to optimize TCO across your choices. It brings unparalleled scale and performance to your mission-critical applications and analytics while securing future readiness for evolving data models.
- **Single pane of glass** visibility for your infrastructure and workloads. Register multi-cloud, including public and private in a single management console and launch virtual analytic workspaces or virtual warehouses within each environment as needed.
- **Secure data and workload migration** to protect your enterprise data and deliver it where is needed. Securely manage data and meta-data migration across all environments.
- **Unified and multi-function Analytics** for cloud-native workloads whether real-time or batch. Integrates data management and analytics experiences across the entire data lifecycle for data anywhere.

-
- **Hybrid and multi-cloud data warehouse** service for all modern, self-service, and advanced analytics use cases, at scale.
 - **Track and Audit everything** across entire ecosystem of CDIP deployments.

CDIP with CDP Private Cloud Hybrid Uses Cases

With the increasing hybrid cloud adoption due to increasing data volume and variety, CDIP addresses uses cases that caters the needs of today's demand of hybrid data platform such as:

- **Hybrid Workload** - Offload workload on-premises to cloud or vice-versa as per the requirements or auto-scale during peak hours due to real-time urgency or seasonality Cloudera Replication Manager and Cloudera Workload Manager.
- **Hybrid Pipelines** - Implement and optimize data pipelines for easier management. Automate and orchestrate your data pipelines as per demand or where it is needed the most. Implement secure data exchange between choice of your cloud and on-premises data hub at scale.
- **Hybrid Data Integration** - Integrate data sources among clouds. Simplify application development or ML model training that needs on-premises data sources or cloud-native data stores.
- **Hybrid DevOps** - Accelerate development with dev sandboxes in the cloud, however, production runs on-premises.
- **Hybrid Data Applications** - Build applications that runs anywhere for cost, performance, and data residency.

Cisco Data Intelligence Platform with Cloudera Data Platform

Cisco developed numerous industry leading Cisco Validated Designs (reference architectures) in the area of Big Data, compute farm with Kubernetes (CVD with RedHat OpenShift Container Platform) and Object store.

A CDIP architecture as a private cloud can be fully enabled by the Cloudera Data Platform with the following components:

- Data lake enabled through CDP PvC Base
- Private Cloud with compute on Kubernetes can be enabled through CDP Private Cloud Data Services and
- Exabyte storage enabled through Apache Ozone

Apache Ozone Brings the Best of HDFS and Object Store

- Overcomes HDFS limitations
 - Can support billions of files (unlike HDFS which only supports up to 500 million files).

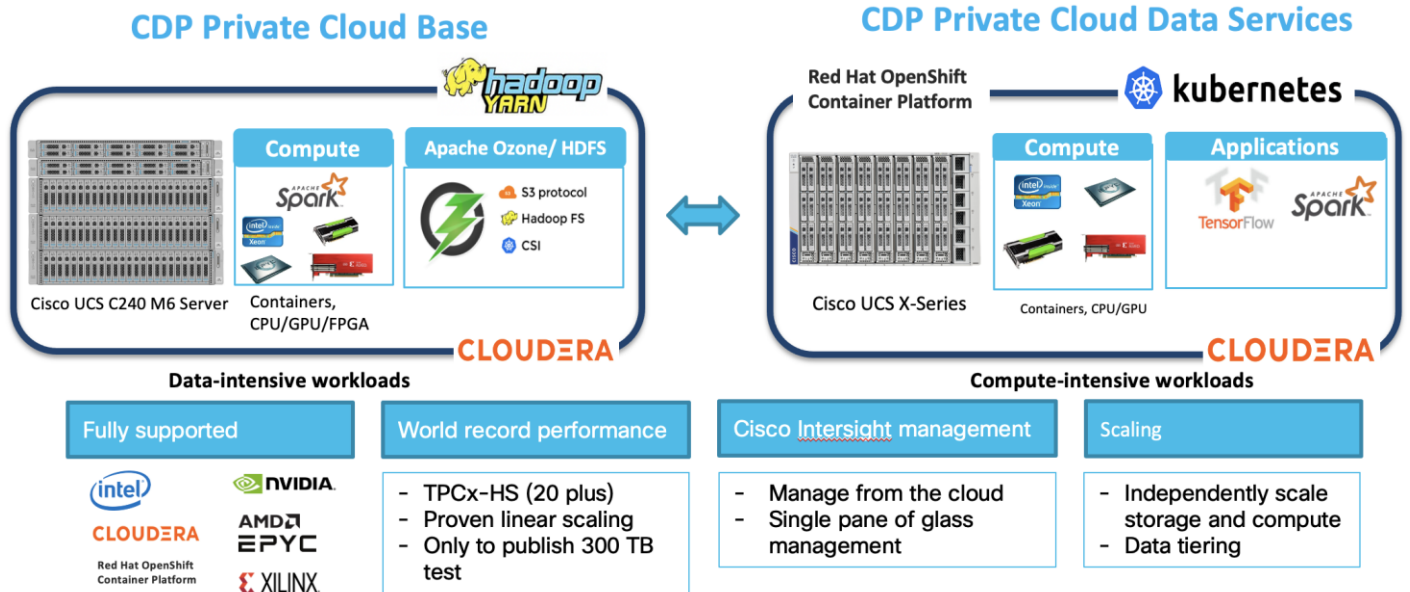
-
- Can currently support 400 TB/ node with potential of supporting 1PB /node at a later point in time unlike HDFS which only supports up to 100 TB/node.
 - Supports 16TB drives unlike HDFS which only supports up to 8TB drives.
 - Exabyte scale.
 - Overcome Object Store limitations
 - Can support large files with linear performance. Like HDFS, Apache Ozone breaks files into smaller chunks (Object stores fail to do this and don't perform linearly with large files), solving the large file problems often hit in object stores.
 - Separates control and data plane enabling high performance. Supports very fast reads out of any of the three replicas.
 - Data from HDFS can be easily migrated to Apache Ozone with familiar tools like distcp. Apache Ozone handles both large and small size files.
 - Supports data locality similar to HDFS and Supports Disaggregation of compute and storage
 - Ozone provides an easy to use monitoring and management console using recon
 - Collects and aggregates metadata from components and present cluster state.
 - Metadata in cluster is disjoint across components
 - Applications like Apache Spark, Hive and YARN, work without any modifications when using Ozone. Ozone comes with a Java client library, S3 protocol support, and a command line interface which makes it easy to use and administer.

Benefits of Apache Ozone to Customers

Apache Ozone brings storage consolidation in a data lake and provides customers with the following:

- Lower Infrastructure cost
- Better TCO and ROI on their investment
- Lower datacenter footprint

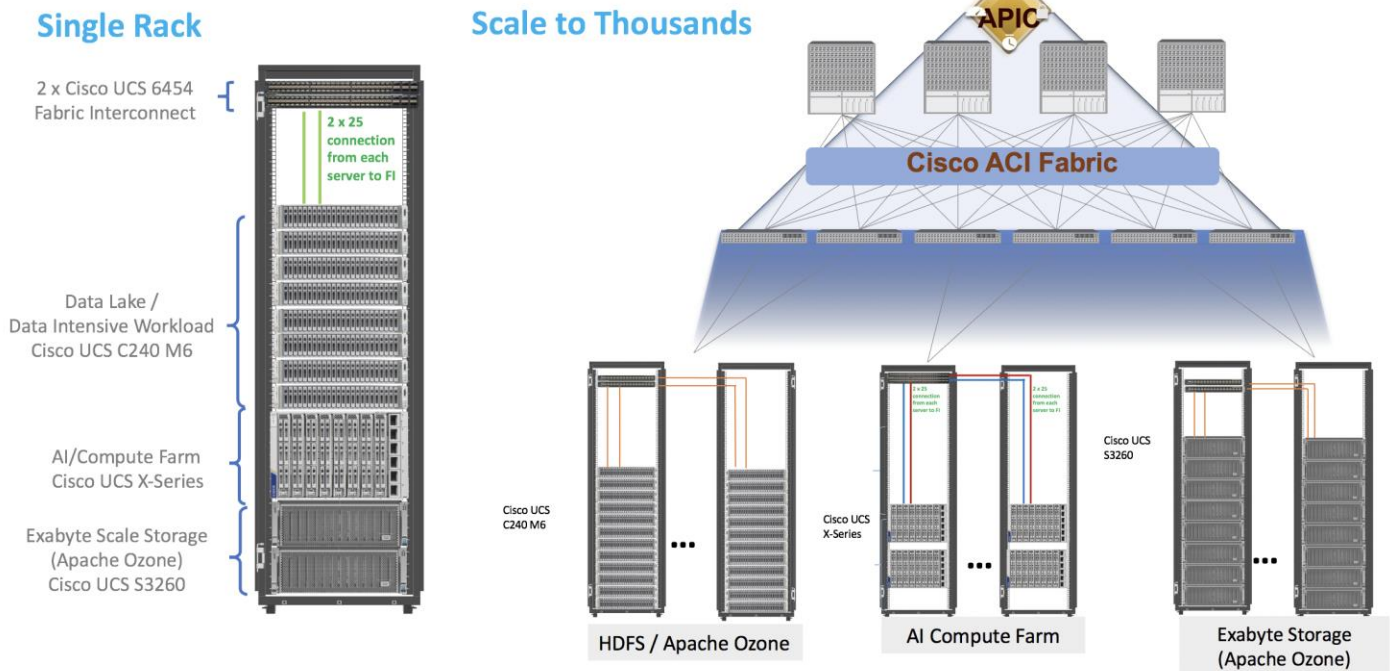
Figure 3. Cisco Data Intelligent Platform with Cloudera Data Platform on Cisco UCS M6



This architecture can start from a single rack and scale to thousands of nodes with a single pane of glass management with Cisco Application Centric Infrastructure (ACI).

Figure 4. Cisco Data Intelligent Platform at Scale

Scaling Cisco Data Intelligence Platform



Reference Architecture

Cisco Data Intelligence Platform reference architectures are carefully designed, optimized, and tested with the leading big data and analytics software distributions to achieve a balance of performance and capacity to address specific application requirements. You can deploy these configurations as is or use them as templates for building custom configurations. You can scale your solution as your workloads demand, including expansion to thousands of servers through the use of Cisco Nexus 9000 Series Switches. The configurations vary in disk capacity, bandwidth, price, and performance characteristics.

Data Lake Reference Architecture

[Table 1](#) lists the data lake, private cloud, and dense storage with Apache Ozone reference architecture for Cisco Data Intelligence Platform.

Table 1. Cisco Data Intelligence Platform with CDP Private Cloud Base (Apache Ozone) Configuration on Cisco UCS M6

	High Performance	Performance	Capacity
Servers	16 x Cisco UCS C240 M6SN Rack Servers with small-form-factor (SFF) drives	16 x Cisco UCS C240 M6 Rack Servers with small-form-factor (SFF) drives	16 x Cisco UCS C240 M6 Rack Servers with large-form-factor (LFF) drives
CPU	2 x 3 rd Gen Intel® Xeon® Scalable Processors 6330 processors (2 x 28 cores, at 2.0 GHz)	2 x 3 rd Gen Intel® Xeon® Scalable Processors 6330 processors (2 x 28 cores, at 2.0 GHz)	2 x 3 rd Gen Intel® Xeon® Scalable Processors 6330 processors (2 x 28 cores, at 2.0 GHz)
Memory	16 x 32 GB RDIMM DRx4 3200 MHz (512 GB)	16 x 32 GB RDIMM DRx4 3200 MHz (512 GB)	16 x 32 GB RDIMM DRx4 3200 MHz (512 GB)
Boot	M.2 with 2 x 960-GB SSDs	M.2 with 2 x 960-GB SSDs	M.2 with 2 x 960-GB SSDs
Storage	24 x 6.4TB 2.5in U2 NVMe and 2 x 3.2TB NVMe	24 x 2.4TB 12G SAS 10K RPM SFF HDD (4K) (or 24 x 7.6TB Enterprise Value 12G SATA SSDs) and 2 x 3.2TB NVMe	16 x 16TB 12G SAS 7.2K RPM LFF HDD(4K) and 2 x 3.2TB NVMe
Virtual interface card (VIC)	4x25G mLOM Cisco UCS VIC 1467	4 x 25G mLOM Cisco UCS VIC 1467	4 x 25G mLOM Cisco UCS VIC 1467
Storage controller	NA	Cisco M6 12G SAS RAID Controller with 4GB FBWC or Cisco M6 12G SAS HBA	Cisco M6 12G SAS RAID Controller with 4GB FBWC or Cisco M6 12G SAS HBA
Network connectivity	Cisco UCS 6454 Fabric Interconnect	Cisco UCS 6454 Fabric Interconnect	Cisco UCS 6454 Fabric Interconnect
GPU (optional)		NVIDIA GPU A10 or NVIDIA GPU A100	NVIDIA GPU A10 or NVIDIA GPU A100


Private Cloud Reference Architecture

[Table 2](#) lists the CDIP private cloud configuration for master and worker nodes.

Table 2. Cisco Data Intelligence Platform with CDP Private Cloud Experiences configuration

	High Core Option
Servers	Cisco UCS X-Series with X210C Blades (Up to 8 Per chassis)
CPU	2 x 3 rd Gen Intel® Xeon® Scalable Processors 6338 processors (2 x 32 cores, at 2.0 GHz)
Memory	16 x 64GB (1TB)
Boot	M.2 with 2 x 960GB SSD
Storage	6 x 3.2TB 2.5in U2 Intel P5600 series NVMe (Ceph [2 drives], Local storage [4 drives])
VIC	4x25 Gigabit Ethernet with Cisco UCS VIC 14425 mLOM
Storage controller	Cisco M6 12G SAS RAID Controller with 4GB FBWC or Cisco M6 12G SAS HBA
Network connectivity	Cisco UCS 6454 Fabric Interconnect

Figure 5. Cisco Data Intelligence Platform on Cisco UCS M6 - Reference Architecture



Component	C240 M6S	C220 M6S	X-Series X210C
CPU	2 x 6330 (28C/2.0GHz)	2 x 6330 (28C/2.0GHz)	2 x 6330 (28C/2.0GHz)
Network	4th Gen FI 6454 / VIC 1467 4 x 25G mLOM	4th Gen FI 6454 / VIC 1467 4 x 25G mLOM	4th Gen FI 6454 / VIC 14425 8 x 25G mLOM
Memory	32G x 16 (512G)	32G x 16 (512G)	32G x 16 (512G)
Drives	6 x 2.4TB 10krpm SFF HDD	4 x 3.8TB NVMe	2 x 1.9TB NVMe
OS Drives	2 x M.2 with 960GB	2 x M.2 with 960GB	2 x M.2 with 960GB

As illustrated in [Figure 6](#), this CVD was designed with the following:

- 24 x Cisco UCS C240 M6 Rack Server running Cloudera Data Platform Private Cloud Base with Apache Ozone
- Cloudera Data Platform Private Cloud Base running the Cloudera manager.

Refer to http://www.cisco.com/go/bigdata_design to build a fully supported CDP Private Cloud Base on CDIP reference architecture. This CVD does not provide the details to build a CDP Private Cloud Base. For detailed instruction, click the following links:


[Cisco Data Intelligence Platform with All NVMe Storage, Cisco Intersight, and Cloudera Data Platform](#)

[Cisco Data Intelligence Platform on Cisco UCS S3260 with Cloudera Data Platform](#)

[Cisco Data Intelligence Platform with Cloudera Data Platform](#)

24 node cluster with Rack#1 hosting 13 x Cisco UCS C240 M6 and Rack#2 hosting 11 x Cisco UCS C240 M6 reference architecture represented [Figure 6](#) representing 25 Gigabit Ethernet link from each of the server connected to a pair of Cisco Fabric Interconnect switches.

Figure 6. Cisco Data Intelligence Platform with Cloudera Data Platform Base running Apache Ozone





The Cisco UCS VIC 1467 provides 10/25Gbps, and the Cisco UCS VIC 1477 provides 40/100Gbps connectivity for the Cisco UCS C-series rack server. For more information see: [Cisco UCS C-Series Servers Managing Network Adapters.](#)



In this solution we configured quad-port mLOM VIC 1467 installed in Cisco UCS C240 M6 server with port 1-2 connected to Fabric A and port 3-4 connected to Fabric B to achieve a total of 50GbE bandwidth per server.



NVMe for Apache Ozone metadata is configured in RAID1 to provide business continuity if there is a hardware failure. This is not needed on compute nodes which only use them for caching. Master Nodes and Data Nodes uses NVMe to store Ozone metadata. The Compute Nodes use NVMe for shuffle and can be in JBOD. The mixed Compute Data Nodes use NVMe for both Apache Ozone metadata and shuffle which requires mount Ozone partitions across both drives as RAID1, with the remaining space used for shuffle/cache as independent JBOD partitions.

Technology Overview

Cisco Data Intelligence Platform

This section describes the components used to build Cisco Data Intelligence Platform, a highly scalable architecture designed to meet a variety of scale-out application demands with seamless data integration and management integration capabilities.

Cisco Data Intelligence Platform powered by Cloudera Data Platform delivers:

- Latest generation of CPUs from Intel (3rd generation Intel Scalable family, with Ice Lake CPUs).
- Cloud scale and fully modular architecture where big data, AI/compute farm, and massive storage tiers work together as a single entity and each CDIP component can also scale independently to address the IT issues in the modern data center.
- World record Hadoop performance both for MapReduce and Spark frameworks published at [TPCx-HS benchmark](#).
- AI compute farm offers different types of AI frameworks and compute types (GPU, CPU, FPGA) to work data for analytics.
- A massive storage tier enables to gradually retire data and quick retrieval when needed on a storage dense sub-systems with a lower \$/TB providing a better TCO.
- Data compression with FPGA, offload compute-heavy compression tasks to FPGA, relieve CPU to perform other tasks, and gain significant performance.
- Seamlessly scale the architecture to thousands of nodes.
- Single pane of glass management with Cisco Intersight.
- ISV Partner ecosystem – Top notch ISV partner ecosystem, offering best of the breed end-to-end validated architectures.
- Pre-validated and fully supported platform.
- Disaggregate Architecture supports separation of storage and compute for a data lake.
- Container Cloud, Kubernetes, compute farm backed by the industry leading container orchestration engine and offers the very first container cloud plugged with data lake and object store.

Containerization

Hadoop 3.0 introduced production-ready Docker container support on YARN with GPU isolation and scheduling. This opened up plethora of opportunities for modern applications, such as micro-services

and distributed applications frameworks comprised of 1000s of containers to execute AI/ML algorithms on peta bytes of data with ease and in a speedy fashion.

Distributed Deep Learning with Apache Submarine

Hadoop community initiated the Apache Submarine project to make distributed deep learning/machine learning applications easily launched, managed, and monitored. These improvements make distributed deep learning/machine learning applications (such as TensorFlow) run on Apache Hadoop YARN, Kubernetes, or just a container service. It enables data scientists to focus on algorithms instead of worrying about underlying infrastructure.

Apache Spark 3.0

Apache Spark 3.0 is a highly anticipated release. To meet this expectation, Spark is no longer limited just to CPU for its workload, it now offers GPU isolation and pooling GPUs from different servers to accelerated compute. To easily manage the deep learning environment, YARN launches the Spark 3.0 applications with GPU. This prepares the other workloads, such as Machine Learning and ETL, to be accelerated by GPU for Spark Workloads. [Cisco Blog on Apache Spark 3.0](#)

Cloudera Data Platform - Private Cloud Base (PvC)

With the merger of Cloudera and Hortonworks, a new “Cloudera” software named Cloudera Data Platform (CDP) combined the best of Hortonwork’s and Cloudera’s technologies to deliver the industry leading first enterprise data cloud. CDP Private Cloud Base is the on-prem version of CDP and CDP Private Cloud Data Services is the on-prem version of Private Cloud to enable compute on Kubernetes with Red Hat OpenShift Container Platform. This unified distribution is a scalable and customizable platform where workloads can be securely provisioned. CDP gives a clear path for extending or refreshing your existing HDP and CDH deployments and set the stage for cloud-native architecture.

CDP Private Cloud Base offers a range of hybrid solutions, including workloads generated with CDP Private Cloud Experiences, where compute activities are isolated from data storage and data may be accessed from data node clusters. This hybrid method manages storage, table structure, authentication, authorization, and governance, laying the groundwork for containerized applications. CDP Private Cloud Base is made up of Apache HDFS, Apache Hive 3, Apache HBase, and Apache Impala, as well as number of additional components for specific applications. You may construct clusters using any combination of these services to meet your business needs and workloads. There are also several pre-configured service bundles for popular workloads.

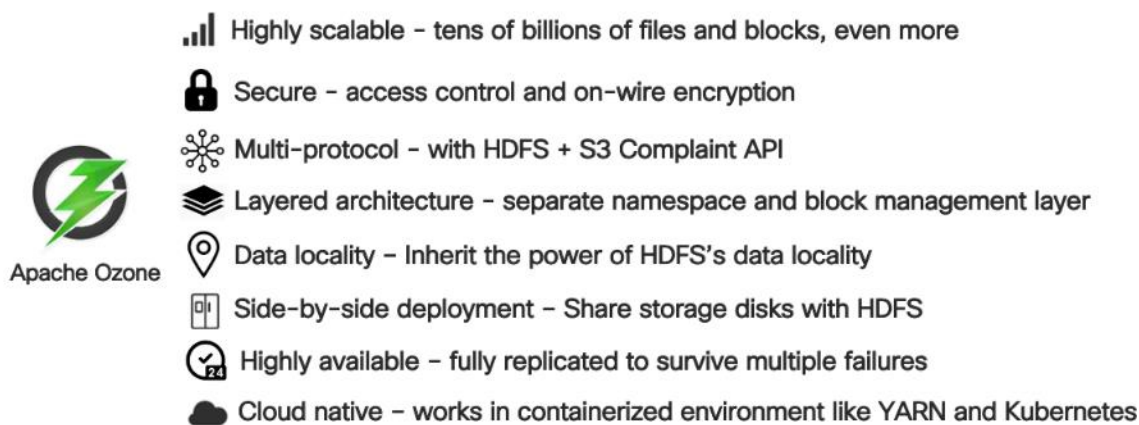
Apache Ozone

Apache Ozone is a distributed, redundant object storage designed for Big Data applications. The major focus of Ozone's architecture is scalability, with the goal of scaling to billions of objects. Ozone isolates namespace and block space management, which makes scaling of distributed systems in Apache Ozone much easier. Apache Ozone can function effectively in containerized environments such as Kubernetes and YARN. Applications using frameworks like Apache Spark, YARN and Hive work natively without any modifications. Ozone is built on a highly available, replicated block storage layer called Hadoop Distributed Data Store (HDDS).

Ozone is a scale-out architecture with minimal operational overheads and long-term maintenance efforts. Ozone can be co-located with HDFS with single security and governance policies for easy data exchange or migration and also offers seamless application portability. Ozone enables separation of compute and storage via the S3 API as well as similar to HDFS, it also supports data locality for applications that choose to use it.

The design of Ozone was guided by the following key principles:

Figure 7. Apache Ozone Key Design Principles



For more details, <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/ozone-overview/topics/ozone-architecture.html>

Cloudera Data Platform Private Cloud Data Services (CDP PvC DS)

Shadow IT can now be eliminated when the CDP Private Cloud is implemented in Cisco Data Intelligence Platform. CDP Private Cloud offers cloud-like experience in customer's on-prem environment. With dis-aggregated compute and storage, complete self-service analytics environment can be implemented, thereby, offering better infrastructure utilization.

Also, CDP Private Cloud offers the persona's Data Scientist, Data Engineer, and Data Analyst, thus providing the right tools to the user improving time-to-value.

Red Hat OpenShift Container Platform (RHOCP) Cluster

Cloudera has selected Red Hat OpenShift as the preferred container platform for CDP Private Cloud. With Red Hat OpenShift, CDP Private Cloud delivers powerful, self-service analytics and enterprise-grade performance with the granular security and governance policies that IT leaders demand.

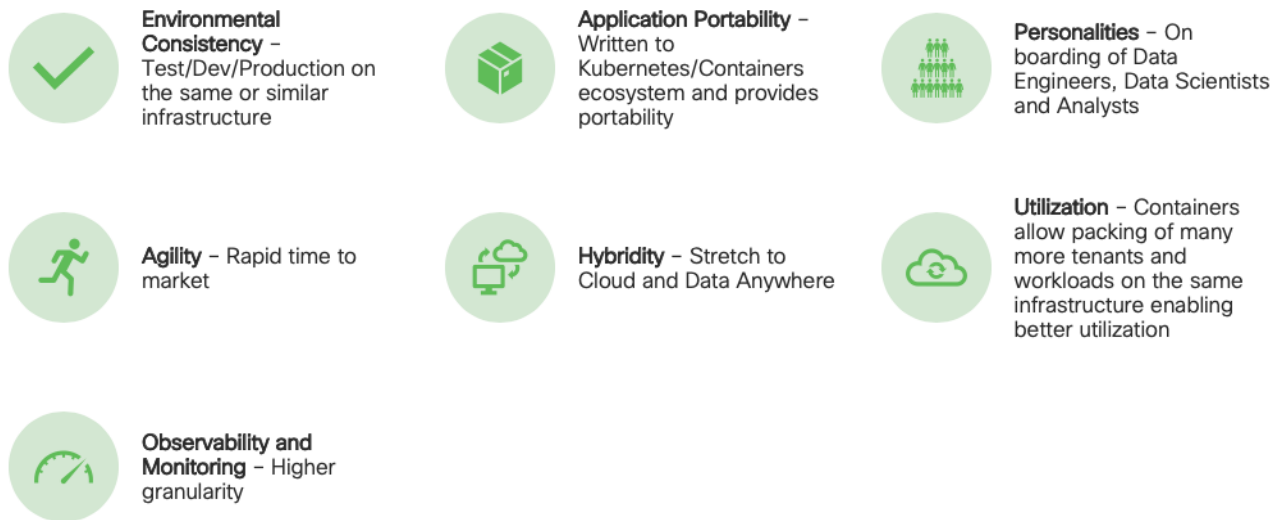
To keep pace in the digital era, businesses must modernize their data strategy for increased agility, ease-of-use, and efficiency. Together, Red Hat OpenShift and CDP Private Cloud help create an essential hybrid, multi-cloud data architecture, enabling teams to rapidly onboard mission-critical applications and run them anywhere, without disrupting existing ones.

Kubernetes

Extracting intelligence from data lake in a timely and speedy fashion is an absolute necessity in finding emerging business opportunities, accelerating time to market efforts, gaining market share, and by all means, increasing overall business agility.

In today's fast-paced digitization, Kubernetes enables enterprises to rapidly deploy new updates and features at scale while maintaining environmental consistency across test/dev/prod. Kubernetes provides the foundation for cloud-native apps which can be packaged in container images and can be ported to diverse platforms. Containers with microservice architecture managed and orchestrated by Kubernetes help organizations embark on a modern development pattern. Moreover, Kubernetes has become in fact, the standard for container orchestration and offers the core for on-prem container cloud for enterprises. It's a single cloud-agnostic infrastructure with a rich open-source ecosystem. It allocates, isolates, and manages resources across many tenants at scale as needed in elastic fashion, thereby, giving efficient infrastructure resource utilization. [Figure 8](#) illustrates how Kubernetes is transforming the use of compute and becoming the standard for running applications.

Figure 8. Compute on Kubernetes is exciting!!!



Spark on Kubernetes

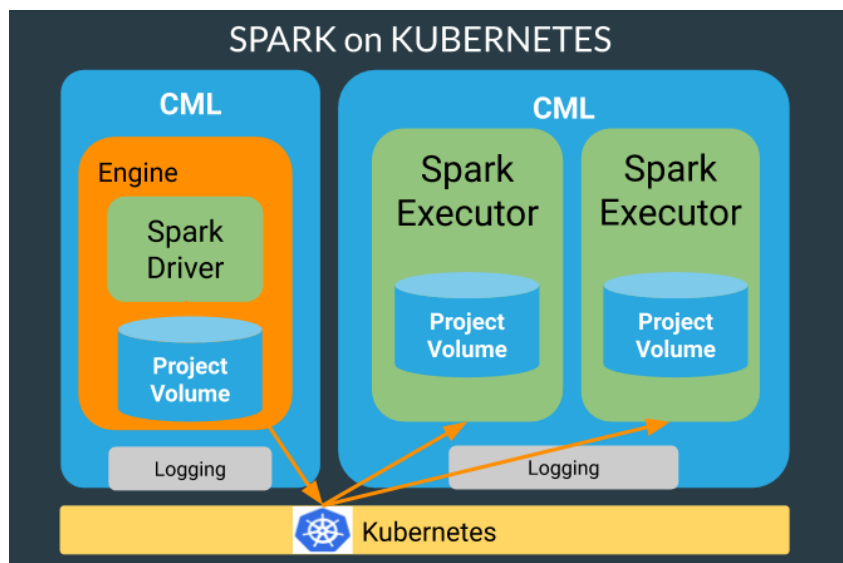
With Spark 2.4.5 along with YARN as a scheduler, comes full support for Apache Spark on Kubernetes as a scheduler. This enables a Kubernetes cluster act as compute layer running Spark workloads for the data lake much of which is used in Cloudera Private Cloud applications.

Spark on Kubernetes has considerably advanced the Hadoop ecosystem, since it made is easier for many public cloud-specific applications and framework use cases to be deployed on-prem; thus, providing hybridity to stretch to cloud anywhere. Kubernetes address gaps that existed in YARN such as lack of isolation and reproducibility and allows workloads to be packaged in docker images. Spark

on Kubernetes also inherit all other in-built features such as auto-scaling, detailed metrics, advanced container networking, security, and so on.

In Cloudera Machine Learning (CML), Spark and its dependencies are bundled directly into the CML engine Docker image. CML supports fully containerized execution of Spark workloads via Spark's support for the Kubernetes cluster backend. Users can interact with Spark both interactively and in batch mode.

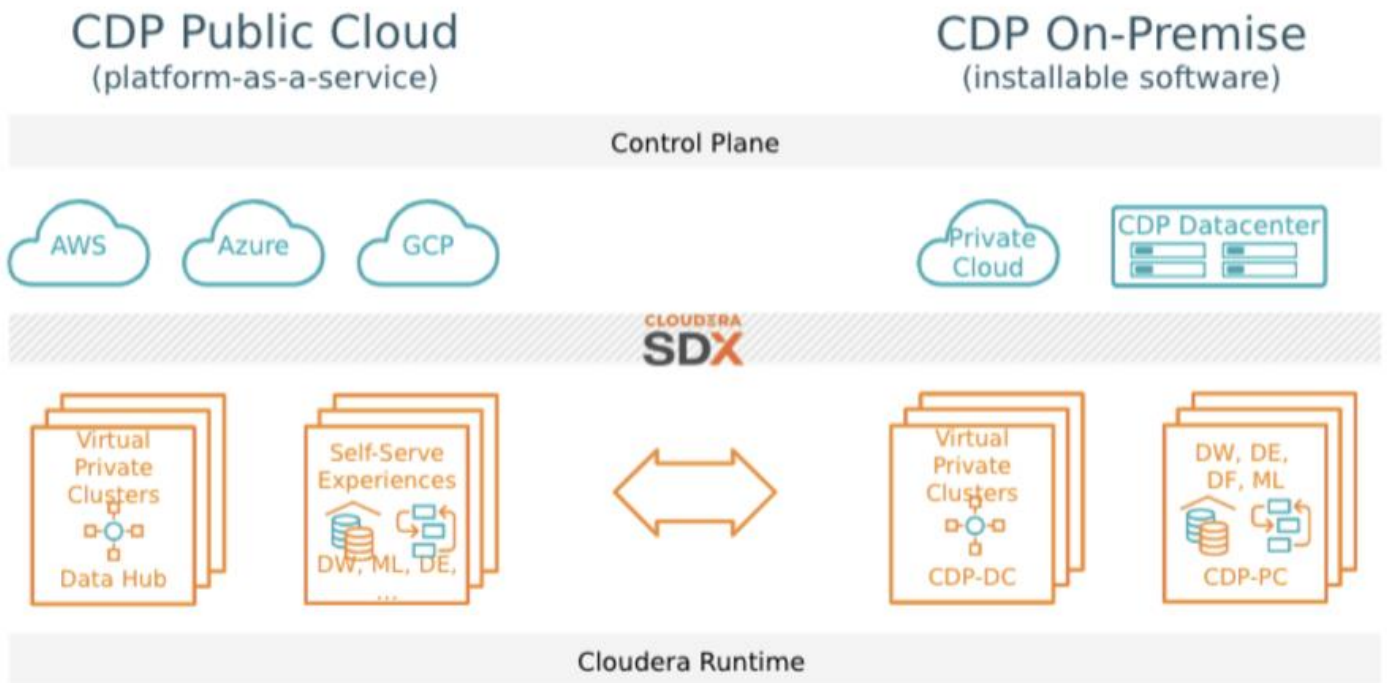
Figure 9. Spark on Kubernetes



Hybrid Architecture

Red Hat OpenShift, being the preferred container cloud platform for CDP private cloud and so is for CDIP, is the market leading Kubernetes powered container platform. This combination is the first enterprise data cloud with a powerful hybrid architecture that decouples compute and storage for greater agility, ease-of-use, and more efficient use of private and multi-cloud infrastructure resources. With Cloudera's Shared Data Experience (SDX), security and governance policies can be easily and consistently enforced across data and analytics in private as well as multi-cloud deployments. This hybridity will open myriad opportunities for multi-function integration with other frameworks such as streaming data, batch workloads, analytics, data pipelining/engineering, and machine learning.

Figure 10. Cloudera Data Platform (CDP) - Hybrid Architecture



Cloud Native Architecture for Data Lake and AI

Cisco Data Intelligence Platform with CDP private cloud accelerates the process of becoming cloud-native for your data lake and AI/ML workloads. By leveraging Kubernetes powered container cloud, enterprises can now quickly break the silos in monolithic application frameworks and embrace a continuous innovation of micro-services architecture with CI/CD approach. With cloud-native ecosystem, enterprises can build scalable and elastic modern applications that extends the boundaries from private cloud to hybrid.

Cisco Unified Computing System

Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce Total Cost of Ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Cisco Intersight

Cisco Intersight is Cisco's systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings

in Total Cost of Ownership (TCO) and to deliver applications faster, so they can support new business initiatives.

Cisco Intersight is a Software as a Service (SaaS) infrastructure management which provides a single pane of glass management of CDIP infrastructure in the data center. Cisco Intersight scales easily, and frequent updates are implemented without impact to operations. Cisco Intersight Essentials enables customers to centralize configuration management through a unified policy engine, determine compliance with the Cisco UCS Hardware Compatibility List (HCL), and initiate firmware updates. Enhanced capabilities and tight integration with Cisco TAC enables more efficient support. Cisco Intersight automates uploading files to speed troubleshooting. The Intersight recommendation engine provides actionable intelligence for IT operations management. The insights are driven by expert systems and best practices from Cisco.

Cisco Intersight offers flexible deployment either as Software as a Service (SaaS) on Intersight.com or running on your premises with the Cisco Intersight virtual appliance. The virtual appliance provides users with the benefits of Cisco Intersight while allowing more flexibility for those with additional data locality and security requirements.

Figure 11. Cisco Intersight



Cisco Intersight has the following:

- Connected TAC
- Security Advisories
- Hardware Compatibility List (HCL) and much more

To learn more about all the features of Intersight, go to:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html>

Cisco UCS 6400 Series Fabric Interconnect

The Cisco UCS 6400 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6400 Series offer line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. ([Figure 12](#) and [Figure 13](#)).


The Cisco UCS 6454 54-Port Fabric Interconnect ([Figure 12](#)) is a One-Rack-Unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE, and Fibre Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has 28 10/25-Gbps Ethernet ports, 4 1/10/25-Gbps Ethernet ports, 6 40/100-Gbps Ethernet uplink ports, and 16 unified ports that can support 10/25-Gbps Ethernet ports or 8/16/32-Gbps Fibre Channel ports. All Ethernet ports are capable of supporting FCoE.

Figure 12. Cisco UCS 6454 Fabric Interconnect



The Cisco UCS 64108 Fabric Interconnect ([Figure 13](#)) is a 2-RU top-of-rack switch that mounts in a standard 19-inch rack such as the Cisco R Series rack. The 64108 is a 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 7.42 Tbps throughput and up to 108 ports. The switch has 16 unified ports (port numbers 1-16) that can support 10/25-Gbps SFP28 Ethernet ports or 8/16/32-Gbps Fibre Channel ports, 72 10/25-Gbps Ethernet SFP28 ports (port numbers 17-88), 8 1/10/25-Gbps Ethernet SFP28 ports (port numbers 89-96), and 12 40/100-Gbps Ethernet QSFP28 uplink ports (port numbers 97-108). All Ethernet ports are capable of supporting FCoE.

Figure 13. Cisco UCS 64108 Fabric Interconnect



Cisco UCS C-Series Rack-Mount Servers

Cisco UCS C-Series Rack-Mount Servers keep pace with Intel Xeon processor innovation by offering the latest processors with increased processor frequency and improved security and availability features. With the increased performance provided by the Intel Xeon Scalable Family Processors, Cisco UCS C-Series servers offer an improved price-to-performance ratio. They also extend Cisco UCS innovations to an industry-standard rack-mount form factor, including a standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology.

It is designed to operate both in standalone environments and as part of Cisco UCS managed configuration, these servers enable organizations to deploy systems incrementally—using as many or as few servers as needed—on a schedule that best meets the organization’s timing and budget. Cisco UCS C-

Series servers offer investment protection through the capability to deploy them either as standalone servers or as part of Cisco UCS. One compelling reason that many organizations prefer rack-mount servers is the wide range of I/O options available in the form of PCIe adapters. Cisco UCS C-Series servers support a broad range of I/O options, including interfaces supported by Cisco and adapters from third parties.

Cisco UCS C240 M6 Rack-Mount Server


The Cisco UCS C240 M6 Rack Server is well-suited for a wide range of storage and I/O-intensive applications such as big data analytics, databases, collaboration, virtualization, consolidation, and high-performance computing in its two-socket, 2RU form factor.

The Cisco UCS C240 M6 Server extends the capabilities of the Cisco UCS rack server portfolio with 3rd Gen Intel Xeon Scalable Processors supporting more than 43 percent more cores per socket and 33 percent more memory when compared with the previous generation.

You can deploy the Cisco UCS C-Series rack servers as standalone servers or as part of the Cisco Unified Computing System managed by Cisco Intersight, or Intersight Managed Mode to take advantage of Cisco® standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility.

These improvements deliver significant performance and efficiency gains that will improve your application performance. The Cisco UCS C240 M6 Rack Server delivers outstanding levels of expandability and performance.


Figure 14. Cisco UCS C240 M6



The Cisco UCS C220 M6 Rack Server is the most versatile general-purpose infrastructure and application server in the industry. This high-density, 1RU, 2-socket rack server delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. You can deploy the Cisco UCS C-Series Rack Servers as standalone servers or as part of the Cisco Unified Computing System managed by Cisco Intersight, Cisco UCS Manager, or Intersight Managed Mode to take advantage of Cisco® standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility.

The Cisco UCS C220 M6 Rack Server extends the capabilities of the Cisco UCS rack server portfolio. The Cisco UCS C220 M6 Rack Server delivers outstanding levels of expandability and performance.

Figure 15. Cisco UCS C220 M6




Cisco UCS X-Series Modular System

The Cisco UCS® X-Series with Cisco Intersight is a modular system managed from the cloud. It is designed to be shaped to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design.

Designed to be managed exclusively from the cloud:

- Simplify with cloud-operated infrastructure
- Simplify with an adaptable system designed for modern applications
- Simplify with a system engineered for the future

Figure 16. Cisco UCS X9508 Modular Chassis with Cisco UCS X210c M6 Compute Node



For more details visit: <https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-x-series-modular-system/index.html>

Simplify Your Data Center

Since we first delivered the Cisco Unified Computing System™ (Cisco UCS) in 2009, our goal has been to simplify your data center. We pulled management out of servers and into the network. We simplified multiple networks into a single unified fabric. And we eliminated network layers in favor of a

flat topology wrapped up into a single unified system. With the Cisco UCS X-Series Modular System we take that simplicity to the next level:

- **Simplify with cloud-operated infrastructure:** Management is moved from the network into the cloud so that you can respond at the speed and scale of your business and manage all of your infrastructure. Shape Cisco UCS X-Series Modular System resources to workload requirements with the Cisco Intersight cloud operations platform. Integrate third-party devices including storage from NetApp, Pure Storage, and Hitachi. Gain intelligent visualization, optimization, and orchestration for all of your applications and infrastructure. Automation drives agility and consistency, helping you reduce time to market while lowering cost and risk.
- **Simplify with an adaptable system designed for modern applications:** Today's cloud-native, hybrid applications are inherently unpredictable. They get deployed and redeployed as part of an iterative DevOps practice. Requirements change often and you need a system that doesn't lock you into one set of resources when you find that you need another. For hybrid applications, and a range of traditional data center applications (see sidebar), you can consolidate onto a single platform that combines the density and efficiency of blade servers with the expandability of rack servers. The result: better performance, automation, and efficiency.
- **Simplify with a system engineered for the future:** Embrace emerging technology and reduce risk with a modular system designed to support future generations of processors, storage, nonvolatile memory, accelerators, and interconnects. Gone is the need to purchase, configure, maintain, power, and cool discrete management modules and servers. Cloud-based management is kept up to date automatically with a constant stream of new capabilities delivered by the Intersight software-as-a-service model.

Support a Broader Range of Workloads

A single server type supporting a broader range of workloads means fewer different products to support, reduced training costs, and increased flexibility. The system supports workloads including the following:

- Virtualized workloads
- Private cloud
- Enterprise applications
- Database management systems
- Infrastructure applications
- Cloud-native applications
- In-memory databases
- Big data clusters

- GPU-accelerated AI/ML workloads

Ready for a Hybrid Cloud World

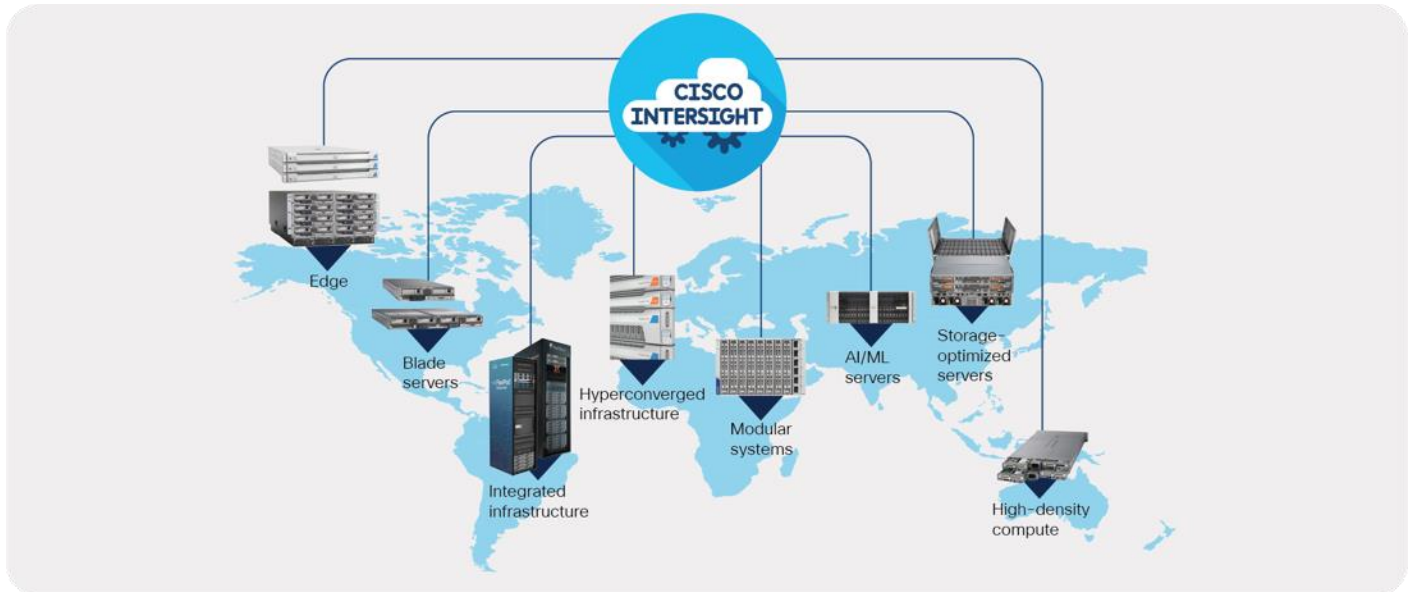
The Cisco Intersight cloud operations platform is the force that transforms the Cisco UCS X-Series Modular System from a set of components into a flexible server platform to propel your most important workloads.

The Cisco UCS X-Series with Intersight is built with a common purpose: to make hardware think like software so that you can easily adapt to a rapidly changing world. Through server profiles, Intersight defines the identity, connectivity, and I/O configuration of your servers and automates the entire infrastructure lifecycle. It's easy to imagine how, as more features are released, the modular system supports a pool of I/O resources: banks of nonvolatile memory, GPU accelerators, specialized ASICs, and massive amounts of NVMe storage. Just as the chassis and Cisco UCS X-Fabric technology are designed to incorporate a constant flow of new capabilities, Cisco Intersight is designed to automatically integrate those technologies into servers along with a constant flow of new, higher-level management capabilities. Software as a service (SaaS) meets modular, infrastructure as code, and the line between hardware and software dissolves.

In its [FutureScape: Worldwide IT Industry 2020 Predictions report](#), IDC predicts that, by 2023, 300 percent more applications will run in the data center and edge locations, 500 million digital applications and services will be developed using cloud-native approaches, and more than 40 percent of new enterprise IT infrastructure will be deployed at the edge. This means that you need a consistent operational approach for all of your infrastructure, wherever it is deployed. With Intersight and the Cisco UCS X-Series you can:

- Define desired system configurations based on policies that use pools of resources provided by the Cisco UCS X-Series. Let Cisco Intersight assemble the components and set up everything from firmware levels to which I/O devices are connected. Infrastructure is code, so your IT organization can use the Intersight GUI, and your DevOps teams can use the Intersight API, the Intersight Service for HashiCorp Terraform, or the many API bindings from languages such as Python and PowerShell.
- Deploy from the cloud to any location. Anywhere the cloud reaches, Intersight can automate your IT processes. We take the guesswork out of implementing new services with a curated set of services we bundle with the Intersight Kubernetes Service, for example.
- Visualize the interdependencies between software components and how they use the infrastructure that supports them with Intersight Workload Optimizer.

Figure 17. Cisco Intersight reaches from the cloud to all of your infrastructure, regardless of its location.



- Optimize your workload by analyzing runtime performance and make resource adjustments and workload placements to keep response time within a desired range. If your first attempt at matching resources to workloads doesn't deliver the results you need, you can reshape the system quickly and easily. Cisco Intersight facilitates deploying workloads into your private cloud and into the public cloud. Now one framework bridges your core, cloud, and edge infrastructure, managing infrastructure and workloads wherever they are deployed.
- Maintain your infrastructure with a consolidated dashboard of infrastructure components regardless of location. Ongoing telemetry and analytics give early detection of possible failures. Reduce risk of configuration drift and inconsistent configurations through automation with global policy enforcement.
- Support your infrastructure with AI-driven root-cause analysis and automated case support for the always-connected Cisco Technical Assistance Center (Cisco TAC). Intersight watches over you when you update your solution stack, helping to prevent incompatible hardware, firmware, operating system, and hypervisor configurations.

Modular Management Architecture

Cisco Intersight is a unified, secure, modular platform that consists of a set of services that bridge applications and infrastructure to meet your specific needs, including:

- Intersight Infrastructure Service
Manage your infrastructure lifecycle, including Cisco data center products, Cisco converged infrastructure solutions, and third-party endpoints.
- Intersight Workload Optimizer

Revolutionize how you manage application resources across any environment with real-time, full-stack visibility to help ensure performance and better cost control.

- Intersight Kubernetes Service

Simplify Kubernetes with automated lifecycle management across your multicloud environment.

- Intersight Virtualization Service

Deploy and manage virtual machines on premises or in the cloud.

- Intersight Cloud Orchestrator

Standardize application lifecycle management across multiple clouds.

Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS VIC 1467

The Cisco UCS VIC 1467 ([Figure 18](#)) is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M6 generation of Cisco UCS C-Series Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBA.

Figure 18. Cisco UCS VIC 1467



Cisco Intersight

Cisco Intersight is Cisco's systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings

in Total Cost of Ownership (TCO) and to deliver applications faster, so they can support new business initiatives.

Cisco Intersight is a Software as a Service (SaaS) infrastructure management which provides a single pane of glass management of CDIP infrastructure in the data center. Cisco Intersight scales easily, and frequent updates are implemented without impact to operations. Cisco Intersight Essentials enables customers to centralize configuration management through a unified policy engine, determine compliance with the Cisco UCS Hardware Compatibility List (HCL), and initiate firmware updates. Enhanced capabilities and tight integration with Cisco TAC enables more efficient support. Cisco Intersight automates uploading files to speed troubleshooting. The Intersight recommendation engine provides actionable intelligence for IT operations management. The insights are driven by expert systems and best practices from Cisco.

Cisco Intersight offers flexible deployment either as Software as a Service (SaaS) on Intersight.com or running on your premises with the Cisco Intersight virtual appliance. The virtual appliance provides users with the benefits of Cisco Intersight while allowing more flexibility for those with additional data locality and security requirements.

Cisco Intersight provides the following features for ease of operations and administration for the IT staff:

- Connected TAC
- Security Advisories
- Hardware Compatibility List (HCL)

To learn more about all the features of Cisco Intersight, go to:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html>

Connected TAC

Connected TAC is an automated transmission of technical support files to the Cisco Technical Assistance Center (TAC) for accelerated troubleshooting.

Cisco Intersight enables Cisco TAC to automatically generate and upload Tech Support Diagnostic files when a Service Request is opened. If you have devices that are connected to Intersight but not claimed, Cisco TAC can only check the connection status and will not be permitted to generate Tech Support files. When enabled, this feature works in conjunction with the Smart Call Home service and with an appropriate service contract. Devices that are configured with Smart Call Home and claimed in Intersight can use Smart Call Home to open a Service Request and have Intersight collect Tech Support diagnostic files.

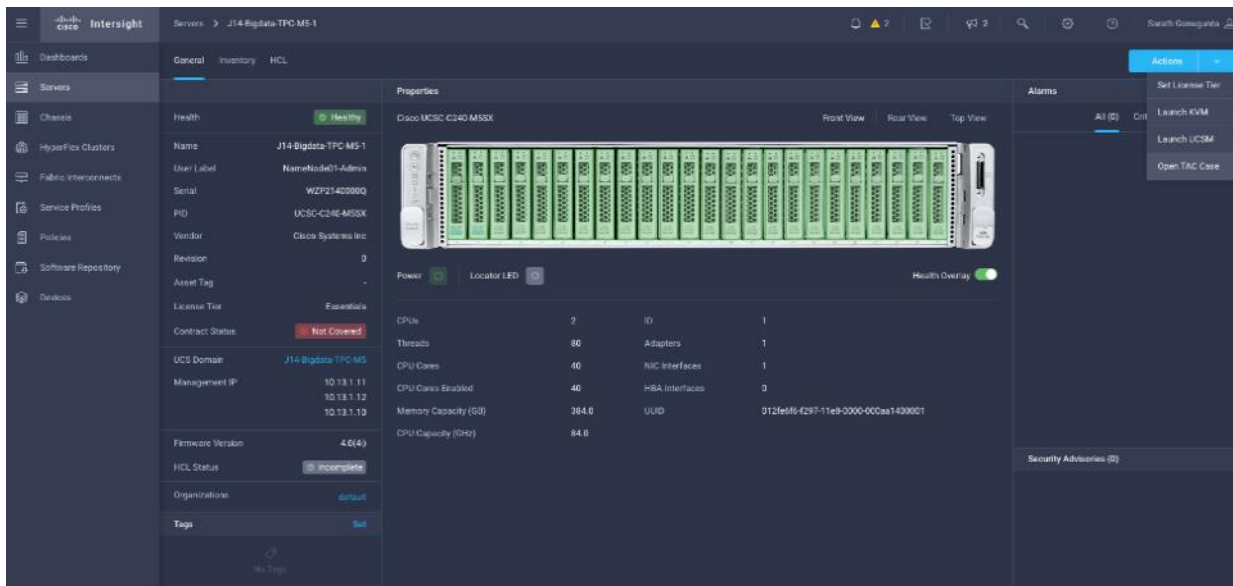
Figure 19. Cisco Intersight: Connected TAC

Cisco Intersight + Cisco TAC + Smart Call Home = Proactive resolution

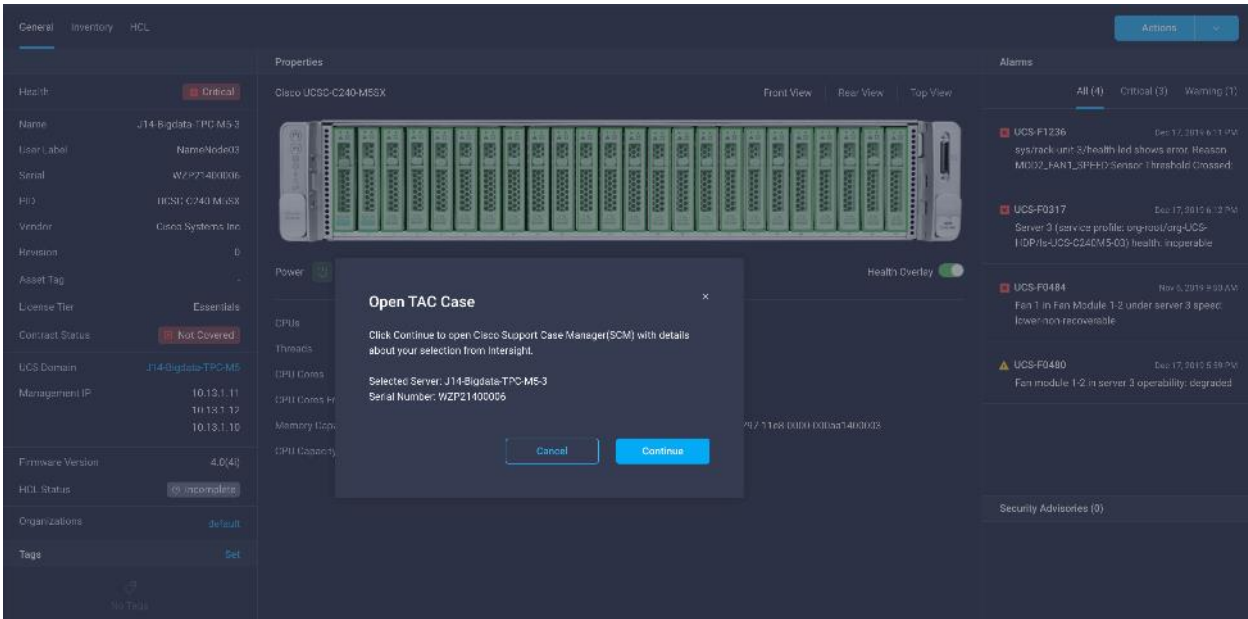


To enable Connected TAC, follow these steps:

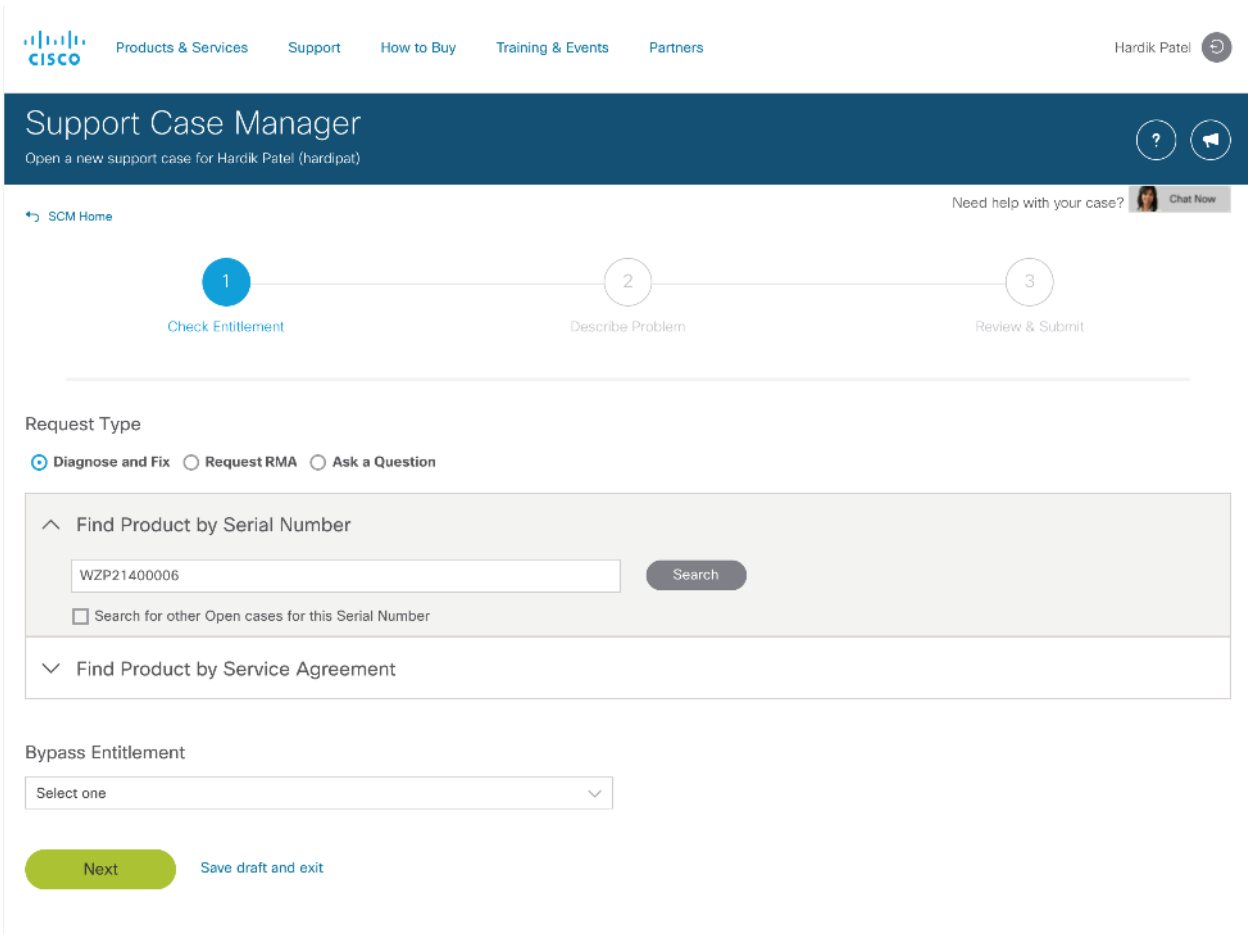
1. Log into [Intersight.com](https://intersight.com).
2. Click the Servers tab. Go to Server > Actions tab. From the drop-down list, click Open TAC Case.
3. Click Open TAC Case to launch the Cisco URL for the support case manager where associated service contracts for Server or Fabric Interconnect is displayed.



4. Click Continue.



5. Follow the procedure to Open TAC Case.



Cisco Intersight Integration for HCL

Cisco Intersight evaluates the compatibility of your Cisco UCS and Cisco HyperFlex systems to check if the hardware and software have been tested and validated by Cisco or Cisco partners. Cisco Intersight reports validation issues after checking the compatibility of the server model, processor, firmware, adapters, operating system, and drivers, and displays the compliance status with the Hardware Compatibility List (HCL).

You can use Cisco UCS Tools, a host utility vSphere Installation Bundle (VIB), or OS Discovery Tool, an open source script to collect OS and driver information to evaluate HCL compliance.

In Cisco Intersight, you can view the HCL compliance status in the dashboard (as a widget), the Servers table view, and the Server details page.



For more information, go to:

https://www.intersight.com/help/features#compliance_with_hardware_compatibility_list_hcl

Figure 20. Example of HCL Status and Driver Recommendation for RHEL 7.8

The screenshot shows the Cisco Intersight interface for a server named 'J14-Bigdata-TPC-M51'. The 'HCL Validation' section shows a status of 'Validated' for 'Server Hardware Compliance'. A modal window titled 'Get Recommended Drivers' is open, showing the following table:

Model	Firmware Vers...	Driver Protocol	Current Driver	Recommended Driver Version
UCSC-RAD...	50.8.0.2649	megaraid_sas	7.705.02.00 r...	60.706.02.00
UCSC-MLOW...	4.3(3b)	enic	2.3.0.53	3.3.21.0.23.756.31
UCSC-MLOW...	4.3(3b)	fcic	-	3.3.2.0.0.1.0.0

Advisories (PSIRTs)

Cisco Intersight sources critical security advisories from the Cisco Security Advisory service to alert users about the endpoint devices that are impacted by the advisories and deferrals. These alerts are displayed as Advisories in Intersight. The Cisco Security Advisory service identifies and monitors and updates the status of the advisories to provide the latest information on the impacted devices, the severity of the advisory, the impacted products, and any available workarounds. If there are no known workarounds, you can open a support case with Cisco TAC for further assistance. A list of the security advisories is shown in Intersight under Advisories.

Figure 21. Intersight Dashboard

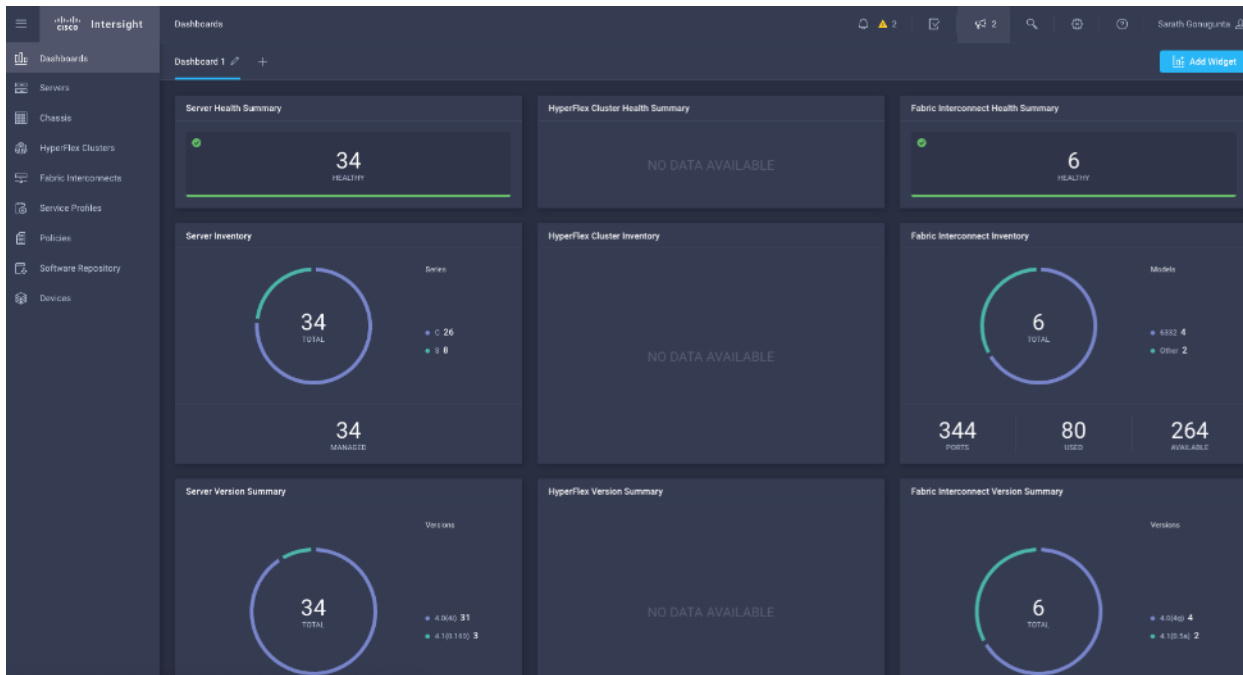
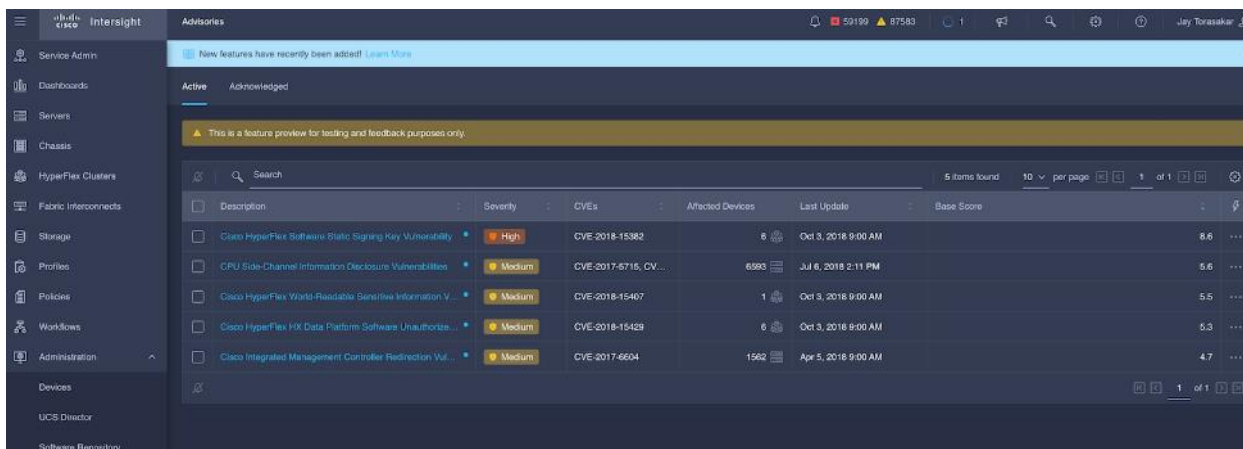


Figure 22. Example: List of PSIRTs Associated with Sample Cisco Intersight Account




The screenshot displays the Cisco Intersight Security Advisories interface. The main content area is titled 'CPU Side-Channel Information Disclosure Vulnerabilities' with a 'Medium' severity rating. The vulnerability ID is 'cisco-cve-2018-0183-sidechannel'. The description explains that these vulnerabilities could allow an attacker to read the bit, bit-operation, and address of memory belonging to other processes or memory allocated to the operating system kernel. It also notes that the first two vulnerabilities (CVE-2017-5753 and CVE-2017-5754) are collectively known as Spectre, and the third (CVE-2017-5754) is known as Meltdown. The 'Affected Devices' section shows '0 affected devices found' in a table with columns for Name, Type, Model / Type, and Firmware / Version. A 'Workarounds/Solutions' section at the bottom states that there are no known current workarounds/solutions for this vulnerability and suggests opening a support case with Cisco TAC.

Cloudera Data Platform (CDP)

CDP is an integrated data platform that is easy to deploy, manage, and use. By simplifying operations, CDP reduces the time to onboard new use cases across the organization. It uses machine learning to intelligently auto scale workloads up and down for more cost-effective use of cloud infrastructure.

Cloudera Data Platform (CDP) Data Center is the on-premises version of Cloudera Data Platform. This new product combines the best of both worlds such as Cloudera Enterprise Data Hub and Hortonworks Data Platform Enterprise along with new features and enhancements across the stack. This unified distribution is a scalable and customizable platform where you can securely run many types of workloads.

Figure 23. Cloudera Data Platform - Unity Release



Cloudera Data Platform provides:

- Unified Distribution: Whether you are coming from CDH or HDP, CDP caters both. It offers richer feature sets and bug fixes with concentrated development and higher velocity.
- Hybrid & On-prem: Hybrid and multi-cloud experience, on-prem it offers best performance, cost, and security. It is designed for data centers with optimal infrastructure.
- Management: It provides consistent management and control points for deployments.
- Consistency: Security and governance policies can be configured once and applied across all data and workloads.
- Portability: Policies stickiness with data, even if it moves across all supported infrastructure.

Cloudera Data Platform Private Cloud Base (CDP PvC Base)

CDP Private Cloud Base is the on-premises version of Cloudera Data Platform. This new product combines the best of Cloudera Enterprise Data Hub and Hortonworks Data Platform Enterprise along with new features and enhancements across the stack. This unified distribution is a scalable and customizable platform where you can securely run many types of workloads.

CDP Private Cloud Base supports a variety of hybrid solutions where compute tasks are separated from data storage and where data can be accessed from remote clusters, including workloads created using CDP Private Cloud Data Services. This hybrid approach provides a foundation for containerized applications by managing storage, table schema, authentication, authorization, and governance.

CDP Private Cloud Base is comprised of a variety of components such as Apache HDFS, Apache Hive 3, Apache HBase, 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. Several pre-configured packages of services are also available for common workloads.

Cloudera Data Platform Private Cloud Data Services (CDP PvC DS)

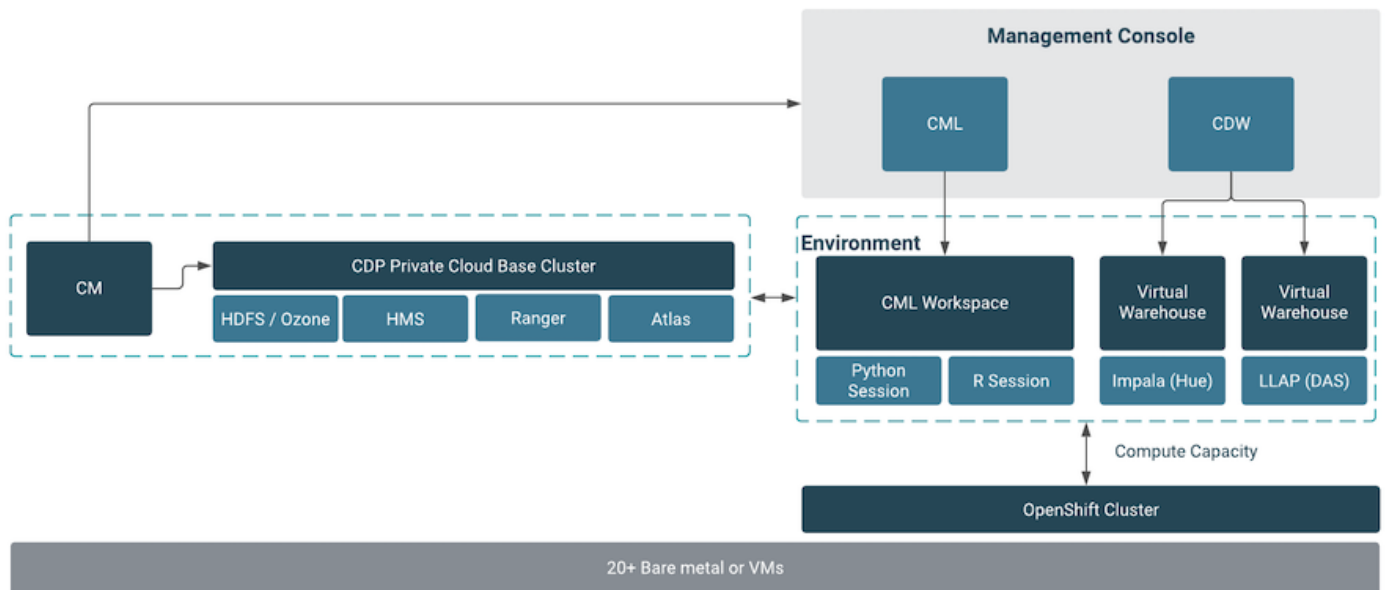
Cloudera Data Platform (CDP) Private Cloud Data Services is the newest on-prem offering of CDP that brings many of the benefits of the public cloud deployments to the on-prem CDP deployments.

CDP Private Cloud Data Services provides a disaggregation of compute and storage and allows independent scaling of compute and storage clusters through the use of containerized applications deployed on Kubernetes. CDP Private Cloud brings both agility and predictable performance to analytic applications. CDP Private Cloud gets unified security, governance, and metadata management through Cloudera Shared Data Experience (SDX), which is available on a CDP Private Cloud Base cluster.

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.

A CDP Private Cloud deployment requires Private Cloud Base cluster and a RedHat OpenShift Kubernetes cluster. The OpenShift cluster is set up on a Bare Metal deployment. 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.

Figure 24. Cloudera Data Platform Private Cloud - Overview



Cloudera Shared Data Experience (SDX)

SDX is a fundamental part of Cloudera Data Platform architecture, unlike other vendors' bolt-on approaches to security and governance. Independent from compute and storage layers, SDX delivers an integrated set of security and governance technologies built on metadata and delivers persistent context across all analytics as well as public and private clouds. Consistent data context simplifies the delivery of data and analytics with a multi-tenant data access model that is defined once and seamlessly applied everywhere.

SDX reduces risk and operational costs by delivering consistent data context across deployments. IT can deploy fully secured and governed data lakes faster, giving more users access to more data, without compromise.

Key benefit and feature of SDX includes:

- **Insightful metadata** - Trusted, reusable data assets and efficient deployments need more than just technical and structural metadata. CDP's Data Catalog provides a single pane of glass to administer and discover all data, profiled and enhanced with rich metadata that includes the operational, social and business context, and turns data into valuable information.
- **Powerful security** - Eliminate business and security risks and ensure compliance by preventing unauthorized access to sensitive or restricted data across the platform with full auditing. SDX enables organizations to establish multi-tenant data access with ease through standardization and seamless enforcement of granular, dynamic, role- and attribute-based security policies on all clouds and data centers.
- **Full encryption** - Enjoy ultimate protection as a fundamental part of your CDP installation. Clusters are deployed and automatically configured to use Kerberos and for encrypted network traffic with Auto-TLS. Data at rest, both on-premises and in the cloud, is protected with enterprise-grade cryptography, supporting best practice tried and tested configurations.
- **Hybrid control** - Meet the ever-changing business needs to balance performance, cost and resilience. Deliver true infrastructure independence. SDX enables it all with the ability to move data, together with its context, as well as workloads between CDP deployments. Platform operational insight into aspects like workload performance deliver intelligent recommendations for optimal resource utilization.
- **Enterprise-grade governance** - Prove compliance and manage the complete data lifecycle from the edge to AI and from ingestion to purge with data management across all analytics and deployments. Identify and manage sensitive data, and effectively address regulatory requirements with unified, platform-wide operations, including data classification, lineage, and modeling.

CDP Private Cloud Management Console

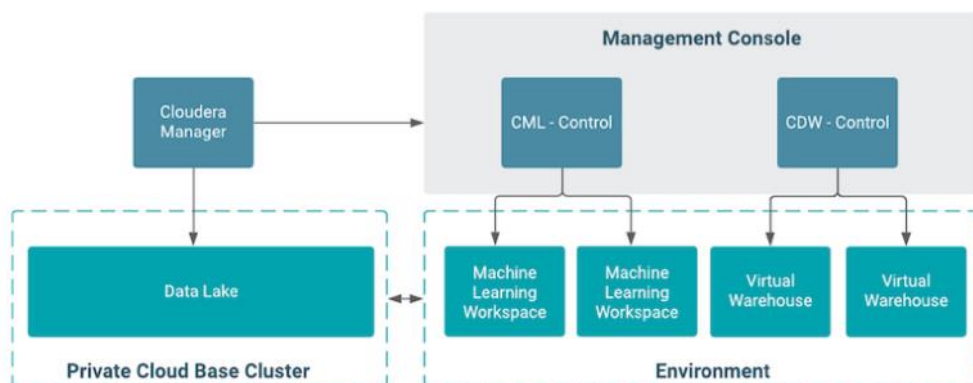
The Management Console is a service used by CDP administrators to manage environments, users, and services.

The Management Console allows you to:

- Enable user access to CDP Private Cloud Data Services, 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.

The following diagram shows a basic architectural overview of the CDP Private Cloud Management Console.

Figure 25. Cloudera Data Platform Private Cloud Management Console



Cloudera Machine Learning (CML)

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.

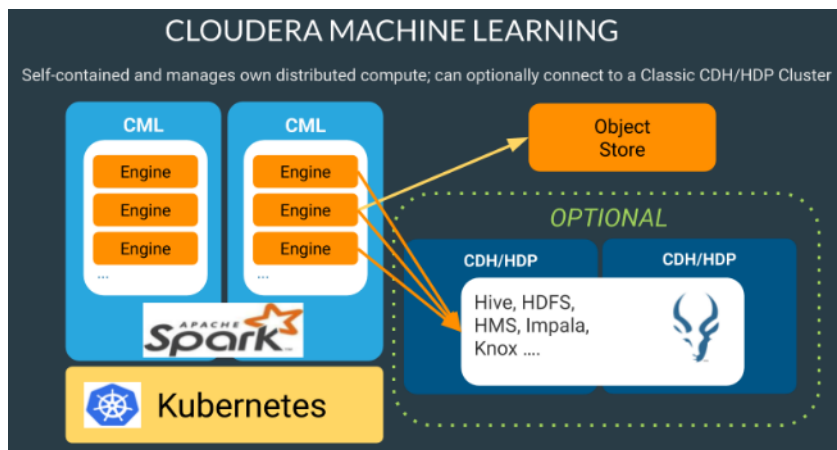
Cloudera Machine Learning (CML) is Cloudera's new cloud-native machine learning service, built for CDP. The CML service provisions clusters, also known as *ML workspaces*, that run natively on Kubernetes.

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*.

Cloudera Machine Learning core capabilities encompasses:

- Seamless portability across private cloud, public cloud, and hybrid cloud powered by Kubernetes
- Rapid cloud provisioning and autoscaling
- Fully containerized workloads - including Python, R, and Spark-on-Kubernetes - for scale-out data engineering and machine learning with seamless distributed dependency management
- High performance deep learning with distributed GPU scheduling and training
- Secure data access across HDFS, cloud object stores, and external databases

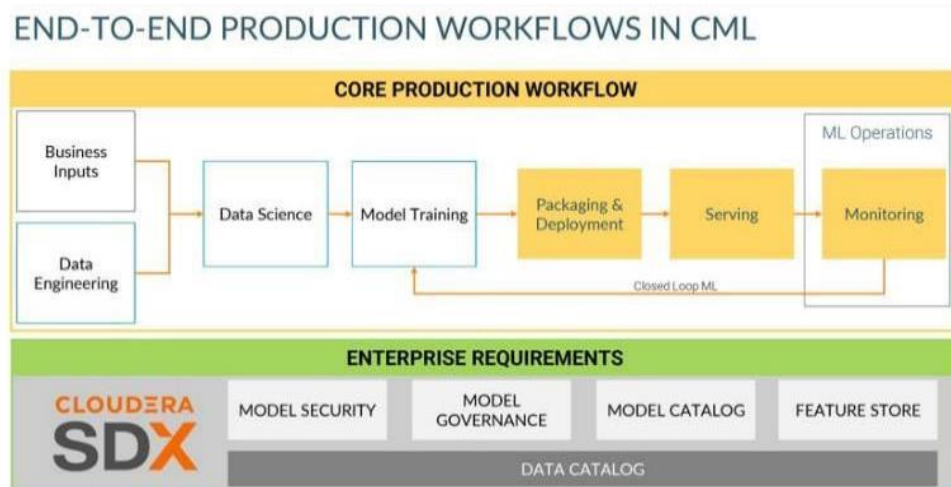
Figure 26. Cloudera Machine Learning



CML MLOps

For enterprises, getting machine learning (ML) models to production and scaling use cases in the business has been a significant challenge. It is estimated that only about 12% of ML models make it to production environments today. To tackle this challenge, Cloudera Machine Learning (CML) MLOps – the most comprehensive and secure production ML platform, built on a 100% open-source standard and fully integrated in CDIP with Cloudera Data Platform. CML breaks the wall to production and enables end-to-end ML workflows at scale.

Figure 27. Cloudera Machine Learning - MLOps



Cloudera Data Warehouse (CDW)

Data Warehouse is a CDP Private Cloud service for self-service creation of independent data warehouses and data marts that autoscale up and down to meet your varying workload demands. The Data Warehouse service provides isolated compute instances for each data warehouse/mart, automatic optimization, and enables you to save costs while meeting SLAs. 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.

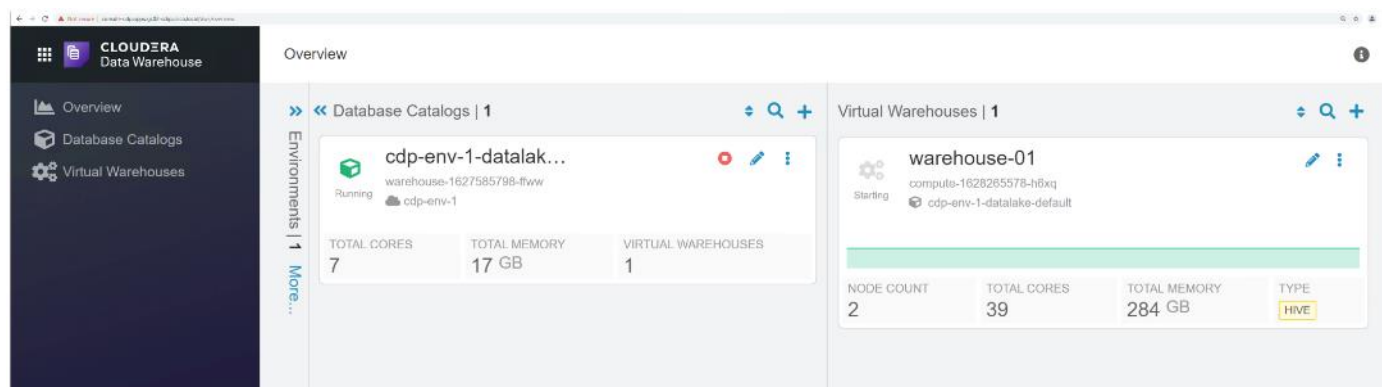
Database Catalogs are Hive MetaStore (HMS) instances and include references to the cloud storage where the data lives. An environment can have multiple Database Catalogs.

The default Database Catalog shares the HMS database with HMS in the base cluster. This enables you to access any objects or data sets created in the base clusters from CDW Virtual Warehouses and vice versa.

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.

Figure 28. Cloudera Data Warehouse (CDW) - Database Catalogs and Virtual Warehouses



Apache Ozone

Apache Ozone is a scalable, redundant, and distributed object store for Hadoop with minimal operational overheads and long-term maintenance efforts. Ozone can be co-located with HDFS with single security and governance policies for easy data exchange or migration and offers 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. Applications using frameworks like Apache Spark, YARN and Hive work natively without any modifications. Ozone is built on a highly available, replicated block storage layer called Hadoop Distributed Data Store (HDDS). Ozone enables separation of compute and storage via the S3 API as well as like HDFS, it also supports data locality for applications that choose to use it.

Ozone consists of volumes, buckets, and keys:

- Volumes are similar to user accounts. Only administrators can create or delete volumes.
- Buckets are similar to directories. A bucket can contain any number of keys, but buckets cannot contain other buckets.
- Keys are similar to files. Each key is part of a bucket, which, in turn, belongs to a volume. Ozone stores data as keys inside these buckets.

When a key is written to Apache Ozone, the associated data is stored on the DataNodes in chunks called blocks. Therefore, each key is associated with one or more blocks. Within the DataNodes, a series of unrelated blocks is stored in a container, allowing many blocks to be managed as a single entity. Apache Ozone is a distributed key-value store that can manage both small and large files alike. While HDFS provides POSIX-like semantics, Ozone looks and behaves like an Object Store.

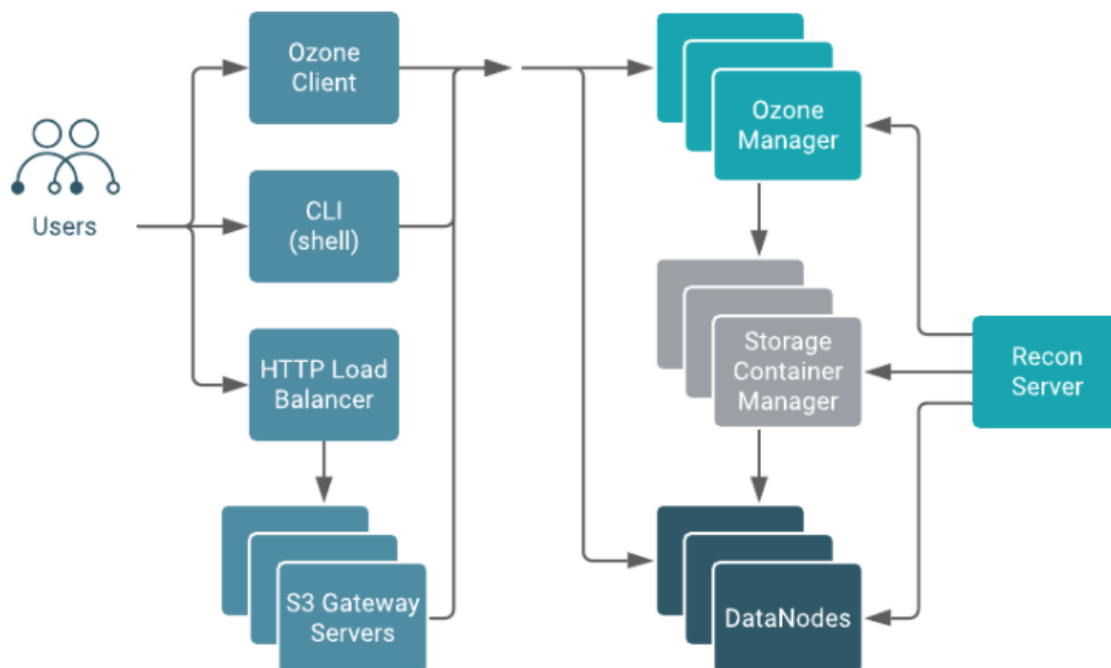
Apache Ozone separates management of namespaces and storage, helping it to scale effectively. Ozone Manager (OM) manages the namespaces while Storage Container Manager (SCM) handles the containers. Ozone Manager (OM) manages the namespace, while Storage Container Manager (SCM) manages the block space (SCM). Volumes, buckets, and keys make up Ozone architecture as shown in [Figure 29](#). In the Ozone world, a volume is comparable to a home directory. It can only be created by an administrator. Buckets are stored using volumes. Users can build as many buckets as they re-

quire once a volume has been created. Data is stored in Ozone as keys that reside within these buckets. The Ozone namespace is made up of several storage volumes. Storage volumes are also used to calculate storage costs.

Apache Ozone components can be summary at the functional levels:

- Metadata data management layer, which is made up of Ozone Manager and Storage Container Manager.
- Data storage layer, which consists of data nodes and is controlled by SCM.
- Ratis' replication layer is used to replicate metadata (OM and SCM) as well as to ensure consistency when data is changed at the data nodes.
- Management server called Recon that communicates with all of Ozone's other components and provides a uniform administration API and UX.
- Protocol bus that permits additional protocols to be added to Ozone. S3 protocol support is presently only available via Protocol Bus. Interface Bus is a general concept that allows you to create new file system or object storage protocols that use the O3 Native protocol.

Figure 29. Apache Ozone Core Components




Apache Ozone brings the following cost savings and benefits due to storage consolidation:

- Lower Infrastructure cost
- Lower software licensing and support cost
- Lower lab footprint
- Newer additional use cases with support for HDFS and S3 and billions of objects supporting both large and small files in a similar fashion.

For more details on Ozone architecture visit, <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/ozone-overview/topics/ozone-introduction.html>

Figure 30. Data Lake Consolidation with Apache Ozone



For more information about Apache Ozone, go to: <https://blog.cloudera.com/apache-ozone-and-dense-data-nodes/>

Persistent Storage for Kubernetes

Workloads deployed in containers and orchestrated via Kubernetes(K8) are either stateless or stateful. By default, K8 workloads are stateless. Stateless applications don't persist, which means it uses temporary storage provided within K8 and destroys once the application or pod is terminated. That's why we call containers are ephemeral in nature, data associated with containers can be lost once the container is terminated or accidentally crashed. Furthermore, data can't be shared among other containers either.

For stateful application, persistent storage is the first "must have" requirement. Kubernetes supports various persistent storage solutions that help addressing this problem and support stateful workloads in a containerized environment. Kubernetes introduces the concept of Persistent Volumes, which exist independently of containers, survive even after containers shut down, and can be requested and consumed by containerized workloads.

There are various methods of providing persistent storage to containers. However, in this reference design, Red Hat OpenShift Container Storage is used to provide persistent volume for Cloudera Private Cloud control plane and Cloudera Machine Learning backed by Red Hat OpenShift Container Platform.

Red Hat OpenShift Container Storage (OCS)

OCS is software-defined storage integrated with and optimized for Red Hat OpenShift Container Platform. OpenShift Container Storage 4.6 is built on Red Hat Ceph® Storage, Rook, and NooBaa to provide container native storage services that support block, file, and object services.

Leveraging the Kubernetes Operator framework, OpenShift Container Storage (OCS) automates a lot of the complexity involved in providing cloud native storage for OpenShift. OCS integrates deeply into cloud native environments by providing a seamless experience for scheduling, lifecycle management, resource management, security, monitoring, and user experience of the storage resources.

To deploy OpenShift Container Storage, the administrator can go to the OpenShift Administrator Console and navigate to the OperatorHub to find the OpenShift Container Storage Operator

OpenShift Container Storage may be used to provide storage for several workloads:

- Block storage for application development and testing environments that include databases, document stores, and messaging systems.
- File storage for CI/CD build environments, web application storage, and for ingest and aggregation of datasets for machine learning.
- Multi-cloud object storage for CI/CD builds artifacts, origin storage, data archives, and pre-trained machine learning models that are ready for serving.

To enable user provisioning of storage, OCS provides storage classes that are ready-to-use when OCS is deployed. OpenShift Container Storage uses the following operators:

-
- The OpenShift Container Storage (OCS) Operator

A meta-operator that codifies and enforces the recommendations and requirements of a supported Red Hat OpenShift Container Storage deployment by drawing on other operators in specific, tested ways. This operator provides the storage cluster resource that wraps resources provided by the Rook-Ceph and NooBaa operators.

- The Rook-Ceph operator

This operator automates the packaging, deployment, management, upgrading, and scaling of persistent storage provided to container applications, and infrastructure services provided to OpenShift Container Platform. It provides the Ceph cluster resource, which manages the pods that host services such as the Object Storage Daemons (OSDs), monitors, and the metadata server for the Ceph file system.

- The NooBaa operator

This operator provides the Multi-cloud Object Gateway, an S3 compatible object store service that allows resource access across multiple cloud environments.

Solution Design

Infrastructure and Software Requirements

This CVD explains the architecture and deployment procedures for Cloudera Data Platform Private Cloud Base on a 24-node cluster using Cisco Data Intelligence Platform for Big Data and Analytics. The solution provides the details to configure CDP PvC Base on the bare metal RHEL infrastructure.


As illustrated in [Figure 31](#), this CVD was designed with the following:





- 3 x Cisco UCS C220 M6 Rack Servers Master Nodes for CDP PvC Base with Apache Ozone
- 3 x Cisco UCS C220 M6 Rack Servers Master Nodes for CDP PvC Base
- 16 x Cisco UCS C240 M6 Rack Servers data Nodes for CDP PvC Base with Apache Ozone
- Cloudera Data Platform Private Base running Apache Ozone on the RedHat Enterprise Linux
- 2 Cisco R42610 standard racks
- 4 Vertical Power distribution units (PDUs) (Country Specific)
- Cloudera Data Platform Private Cloud Base (the data lake) which is not detailed in this CVD but is extensively explained in the CVDs published here: http://www.cisco.com/go/bigdata_design.

Physical Topology

Two-rack consists of total four vertical PDUs and two Cisco UCS Fabric Interconnect with 24 x Cisco UCS C240 M6 Rack Servers connected to each of the vertical PDUs for redundancy. This ensures availability during power source failure. [Figure 31](#) illustrates a 2x25 Gigabit Ethernet link from each server is connected to both Fabric Interconnects. (port 1-2 connected to FI - A and port 3-4 connected to FI - B)

Figure 31. Cisco Data Intelligence Platform with Cloudera Private Cloud Base running Apache Ozone



-  Please contact your Cisco representative for country-specific information.
-  [Intel® Virtual RAID on CPU \(VROC\) with VMD](#) configured RAID 1 for NVMe drives to provide business continuity for ozone metadata in case of hardware failure. This is not needed on pure compute nodes which only use them for caching.
-  The masters and storage dataNodes use NVMe to store Ozone metadata. With each data node hosting 384+ TB data, reducing the chance of a node going offline because of a single NVMe failure. The compute nodes use NVMe for shuffle (Spark, MR, Tez) and for caching (LLAP).
-  The "hybrid" mixed compute dataNodes use NVMe for both ozone metadata and shuffle (spark, mr, tez) + caching (llap). They should mount ozone partitions across both drives as

RAID1 (800GB), with the remaining space used for shuffle/cache as independent JBOD partitions.



Minimum starter configuration is 3 master nodes and 9 dataNodes. This will support erasure coding rs(6,3) in the future. Additional DataNodes can be added in increments of 1 to increase storage.

Logical Topology

Port Configuration on Cisco UCS Fabric Interconnect 6454

[Table 3](#) lists the port configuration on Cisco UCS Fabric Interconnect 6454.

Table 3. Port Configuration on Cisco UCS Fabric Interconnect 6454

Port Type	Port Number
Server	1-48
Uplink Network	49-54

Server Configuration and Cabling for Cisco UCS C240 M6


The Cisco UCS C240 M6 Rack Server is equipped with two 3rd Gen Intel Xeon Scalable Family Processor 6330 (2 x 28 cores, 2.0 GHz), 512 GB of memory (16 x 32GB @ 3200MHz), Cisco UCS Virtual Interface Card 1467, 2 x 3.8TB Intel U.2 P5500 series NVMe, M.2 with 2 x 960GB SSDs for Boot and 24 x 2.4TB 10k rpm 4kn HDD data disks.

[Figure 32](#) illustrates the port connectivity between the Cisco UCS Fabric Interconnect 6454 and Cisco UCS C240 M6 Rack Server with mLOM VIC 1467. Twenty four Cisco UCS C240 M6 servers are installed in this configuration.

For information on physical connectivity and single-wire management, go to:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c-series_integration/ucsm4-2/b_C-Series-Integration_UCSM4-2.html

Figure 32. Network Connectivity for Cisco UCS C240 M5 Rack Server






Software Distributions and Firmware Versions

The software distributions required versions are listed in [Table 4](#).

Table 4. Software Distribution and Version

Layer	Component	Version or Release
Compute	Cisco UCS C240 M6	4.2.1d.0.0730210924
Network	Cisco UCS Fabric Interconnect	4.2(1d)
	Cisco UCS VIC 1467 Firmware	5.2(1)
Storage	Cisco 12G Modular RAID Controller	52.15.0-3988
	LSI MegaRAID SAS Driver	07.715.03.00
	SAS Expander	65160700
Software	Red Hat Enterprise Linux Server (CDP Private Cloud Base)	8.2
	Cloudera Data Platform Private Cloud Base	7.1.7

Layer	Component	Version or Release
	Cloudera Manager	7.4.4
	Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.7.0-551
	Apache Spark	2.4.7.7.1.7.0-551

-
-  The Cisco latest drivers can be downloaded here: <https://software.cisco.com/download/home>.
 -  For Cloudera Private Cloud Base and Experiences versions and supported features, go to: <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/runtime-release-notes/topics/rt-pvc-runtime-component-versions.html>
 -  For Cloudera Private Cloud Base requirements and supported version, go to: <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/installation/topics/cdpdc-requirements-supported-versions.html>
-

Solution Prerequisites

There are many platform dependencies to enable Cloudera Data Platform Private Base running on Cisco Data Intelligence Platform.

The following are the prerequisites needed to enable this solution:

- Network requirements
- Security requirements
- Operating System requirements
- Cloudera requirements

Network Requirements

Network Bandwidth Between HDFS/Ozone and Private Cloud

Cloudera CDP PvC Base cluster that houses HDFS and Apache Ozone storage and Cloudera Private Cloud compute-only clusters should be reachable with no more than a 3:1 oversubscription in order to be able to read from and write to the base cluster. The recommended network architecture is Spine-Leaf between the spine and leaf switches. Additional routing hops should be avoided in production and ideally both HDFS/Ozone storage and Cloudera Private Cloud are on the same network.

For more information, go to: <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/concepts/topics/cm-vpc-networking.html>

NTP

Both Cloudera Base and Cloudera Private Cloud cluster should have their time synched with the NTP Clock time from same the NTP source. Also make sure, Active Directory server where Kerberos is set-up for data lake and for other services must also be synched with same NTP source.

DNS



DNS is required for this solution. It is the requirement for setting up Active Directory, Kerberos, Cloudera Manager authentication, and OpenShift.

DNS must have the following:

- Each host whether Cloudera Manger or Red Hat OpenShift must be accessible via DNS.
- DNS must be configured for forward AND reverse for each host. Reverse is required for Kerberos authentication to the Base Cloudera cluster.
- Cloudera Manager host must be able to resolve hostname of Red Hat OpenShift ingress/route via DNS of wildcard entry to load balancer of OpenShift Container Platform.

-
- Service DNS entry must be configured for ETCD cluster. For each control plane machine, OpenShift Container Platform also requires an SRV DNS record for etcd server on that machine with priority 0, weight 10 and port 2380.
 - A wildcard DNS entry is required for resolving the ingress/route for applications.

Cloudera Data Platform Requirements

Database for Cloudera Manager

Cloudera Manager/Runtime database must be Postgresql 10 or higher. Additional databases are on roadmap. Please review [Cloudera support matrix](#) for the database version supported with Cloudera Manager and CDP PvC Base.



Cloudera Data Warehouse requires Postgresql 10 must be configured with SSL.

Licensing Requirements

The cluster must be setup with a license with entitlements for installing Cloudera Private Cloud Base. Initial trial period of 60 days doesn't require a license file.

JDK 11

The cluster must be configured with JDK 11, JDK8 is not supported. You can use Oracle, OpenJDK 11.04, or higher. JAVA 11 is a JKS requirement and must be met. In this CVD we used Oracle JDK 11.0.9.

Kerberos

Kerberos must be configured using an Active Directory (AD) or MIT KDC. The Kerberos Key Distribution Center (KDC) will use the domain's Active Directory service database as its account database. An Active Directory server is recommended for default Kerberos implementations and will be used in the validation of this solution. Kerberos will be enabled for all services in the cluster.



Red Hat IPA/Identity Management is currently not supported.

Configure Cloudera Manager with TLS

The cluster must be configured with TLS. AutoTLS is recommended which uses an internal certificate authority (CA) created and managed by Cloudera Manager. Auto-TLS automates the creation of an internal certificate authority (CA) and deployment of certificates across all cluster hosts.

Manual TLS add additional requirements and must be met according to the guidelines mentioned here: <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/security-encrypting-data-in-transit/topics/cm-security-tls-comparing.html>

TLS uses JKS-format (Java KeyStore)

Cloudera Manager Server, Cloudera Management Service, and many other CDP services use JKS formatted keystores and certificates. Java 11 is required for JKS.

Required Services

The following minimum services must be configured and setup in Data Lake for private cloud registration process:

- HDFS, Hive Metastore, Ranger, and Atlas.
- There are other dependent services such as Solr for Ranger and HBase and Kafka for Atlas that must be configured.
- All services within the cluster must be in good health. Otherwise, environment registration for private cloud will fail.

Kerberos configuration for Ozone

Ozone depends on Kerberos to make the clusters secure. To enable [Kerberos security in an Ozone cluster](#), you must set the parameters:

```
ozone.security.enabled - true
hadoop.security.authentication - kerberos
```

Ozone Manager in High Availability

Configuring High Availability (HA) for the Ozone Manager (OM) allows you to operate redundant Ozone Managers on your Ozone cluster, preventing the cluster from experiencing a single point of failure in terms of namespace management. In addition, Ozone Manager HA guarantees that read and write activities with client applications are not interrupted.

One leader OM node and two or more follower nodes are used in a High Availability (HA) setup of the Ozone Manager (OM). Client requests are read and written by the leader node services. The follower nodes closely monitor the leader's updates in order for one of the follower nodes to take over the leader's activities in the case of a failure.

Storage Controller Manager in High Availability

Configuring High Availability (HA) for the Storage Container Manager (SCM) in an Ozone cluster to handle the various forms of storage information prevents a single point of failure and guarantees the SCM's continuous interactions with the Ozone Manager (OM) and the DataNodes.

The essential SCM activities, such as managing client requests such as assigning containers, local actions such as deleting pipelines, and processing DataNode changes, are handled differently when an Ozone cluster has Storage Container Manager (SCM) High Availability (HA) set.



Storage controller manager high availability feature enablement requires CDP PvC Base version 7.1.7 or higher.

Configure Security for Ozone

Ranger Configuration for Ozone

In CDP PvC Base cluster, to configure [Apache Ranger service for Ozone](#) to secure access to Ozone data. Go to Cloudera Manager and enable the Ranger service instance. After Ranger service is installed in the cluster.

Go to the Ozone service > Configuration tab > Search for ranger_service.
Select the RANGER service checkbox.
Enter a Reason for Change, and then click Save Changes to save the property changes.
Restart the cluster.



Set up policies in Ranger for the users to have the right access permissions to the various Ozone objects such as buckets and volumes. When using Ranger to provide a particular user with read/write permissions to a specific bucket, you must configure a separate policy for the user to have read access to the volume in addition to policies configured for the bucket.

Summary

Evolving workloads need a highly flexible platform to cater to various requirements, whether data-intensive (data lake) or compute-intensive (AI/ML/DL or container workloads) or just storage-dense (object store). An infrastructure to enable this evolving architecture—one that is able to scale to thousands of nodes—requires strong attention to operational efficiency.

To provide a seamless operation of the application at this scale, you need the following:

- An infrastructure automation with centralized management
- Deep telemetry and simplified granular troubleshooting capabilities
- Multi-tenancy for application workloads, including containers and micro-services, with the right level of security and SLA for each workload

Cisco UCS with Cisco Intersight and Cisco ACI can enable this next-generation cloud-scale architecture, deployed and managed with ease while Cloudera Data Platform Private Cloud Experiences provides the software capabilities to tie these technologies together in a seamless and secure manner.

For More Information

For additional information, see the following resources:

- To find out more about Cisco UCS Big Data solutions, go to: <https://www.cisco.com/go/bigdata>
- To find out more about Cisco UCS Big Data validated designs, go to: https://www.cisco.com/go/bigdata_design
- To find out more about Cisco Data Intelligence Platform, go to: <https://www.cisco.com/c/dam/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/solution-overview-c22-742432.pdf>
- To find out more about Cisco UCS AI/ML solutions, go to: <http://www.cisco.com/go/ai-compute>
- To find out more about Cisco ACI solutions, go to: <http://www.cisco.com/go/aci>
- To find out more about Cisco validated solutions based on Software Defined Storage, go to: <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/software-defined-storage-solutions/index.html>
- Cloudera Data Platform Data Center 7.0 release note, go to: <https://docs.cloudera.com/cdpdc/7.0/release-guide/topics/cdpdc-release-notes-links.html>
- CDP Data Center Requirements and Supported Versions, go to: <https://docs.cloudera.com/cdpdc/7.0/release-guide/topics/cdpdc-requirements-supported-versions.html>

-
- Cloudera supported data storage components, how to: <https://docs.cloudera.com/cdp-private-cloud-base/7.1.7/howto-storage.html>

Bill of Materials

This section provides the BoM for the 8 240M5 Storage Server Hadoop Base Rack. See [Table 5](#) for BOM for the Hadoop Base rack and [Table 6](#) for Red Hat Enterprise Linux License.

Table 5. Bill of Materials for Cisco UCS C240M5SX Nodes Base Rack

Part Number	Description	Qty
UCSC-C240-M6SX	UCS C240 M6 Rack w/o CPU, mem, drives, 2U w 24	24
UCSC-M-V25-04	Cisco UCS VIC 1467 quad port 10/25G SFP28 mLOM	24
CIMC-LATEST	IMC SW (Recommended) latest release for C-Series Servers.	24
UCS-M2-960GB	960GB SATA M.2	48
UCS-M2-HWRAID	Cisco Boot optimized M.2 Raid controller	24
UCSX-TPM-002C	TPM 2.0, TCG, FIPS140-2, CC EAL4+ Certified, for M6 servers	24
UCSC-RAIL-M6	Ball Bearing Rail Kit for C220 & C240 M6 rack servers	24
UCS-DIMM-BLK	UCS DIMM Blanks	384
UCSC-RIS2A-240M6	C240 / C245 M6 Riser2A; (x8;x16;x8);StBkt; (CPU2)	24
UCSC-HSHP-240M6	Heatsink for 2U SFF M6 PCIe SKU	48
UCS-SCAP-M6	M6 SuperCap	24
UCSC-M2EXT-240M6	C240M6 / C245M6 2U M.2 Extender board	24
CBL-SAS12-240M6	C240M6 SAS cable (2U); (Pismo HBA)	24
CBL-RSASR3B-240M6	C240M6 2U x2 Rear SAS/SATA cable; (Riser3B)	24
CBL-SDSAS-240M6	CBL C240M6X (2U24) MB CPU1(NVMe-A) to PISMO BEACH PLUS	24
CBL-R1B-SD-240M6	CBL C240 / C245 M6SX PB+ to Riser 1B	24
CBL-SCAPSD-C240M6	CBL Super Cap for PB+ C240 / C245 M6	24

Part Number	Description	Qty
UCS-CPU-I6330	Intel 6330 2.0GHz/205W 28C/42MB DDR4 2933MHz 28C/42MB DDR4 3200MHz	48
UCS-MR-X32G2RW	32GB RDIMM DRx4 3200 (8Gb)	384
UCSC-RIS1B-240M6	C240 M6 Riser1B; 2xHDD/SSD; StBkt; (CPU1)	24
UCSC-RIS3B-240M6	C240 M6 Riser 3B; 2xHDD; StBkt; (CPU2)	24
UCSC-RAID-M6SD	Cisco M6 12G SAS RAID Controller with 4GB FBWC (28 Drives)	24
UCS-HD24TB10K4KN	2.4 TB 12G SAS 10K RPM SFF HDD (4K)	576
UCS-NVMEI4-I3840	3.8TB 2.5in U.2 Intel P5500 NVMe High Perf Medium Endurance	48
UCSC-PSU1-1600W	Cisco UCS 1600W AC Power Supply for Rack Server	48
CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	48
UCS-SID-INFR-BD	Big Data and Analytics Platform (Hadoop/IoT/ITOA/AI/ML)	24
UCS-SID-WKL-BD	Big Data and Analytics (Hadoop/IoT/ITOA)	24
RACK2-UCS2	Cisco R42612 standard rack, w/side panels	2
CON-SNT-RCK2UCS2	SNTC 8X5XNBD, Cisco R42612 standard rack, w side panels	2
UCS-FI-6454-U	UCS Fabric Interconnect 6454	2
CON-OSP-SFI6454U	SNTC-24X7X4OS UCS Fabric Interconnect 6454	2
UCS-PSU-6332-AC	UCS 6332/ 6454 Power Supply/100-240VAC	4
CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	4
UCS-ACC-6332	UCS 6332/ 6454 Chassis Accessory Kit	2
UCS-FAN-6332	UCS 6332/ 6454 Fan Module	8

Table 6. Red Hat Enterprise Linux License

Part Number	Description	Qty
RHEL-2S2V-3A	Red Hat Enterprise Linux	24
CON-ISV1-EL2S2V3A	3-year Support for Red Hat Enterprise Linux	24

About the Authors

Hardik Patel, Technical Marketing Engineer, Computing Systems Product Group, Cisco Systems, Inc.

Hardik Patel is a Technical Marketing Engineer in Cloud and Compute Product Group. He is currently responsible for architect, design, and development of Cisco Data Intelligence Platform based Big Data infrastructure solutions and performance for private/hybrid cloud and AI/ML. Hardik holds Master of Science degree in Computer Science with various career-oriented certification in virtualization, network, and Microsoft.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Karthik Krishna, Cisco Systems, Inc.
- Silesh Bijjahalli, Cisco Systems, Inc.
- Ali Bajwa, Cloudera
- Karthik Krishnamoorthy, Cloudera
- Arpit Agarwal, Cloudera

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](#) at <https://cs.co/en-cvds>.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)