



Cisco Data Intelligence Platform (CDIP) Powered by AMD

Modernizing your data lake to meet the evolving IT landscape

In today's environment, a voluminous amount of data, reaching exabytes in scale, ends up being stored in data ecosystems. Enterprises are constantly evaluating new sets of data management for processing, transforming, and analyzing, these large amounts of data leading to newer data pipelines evolving beyond the standard data lake.

The rapid advancement of artificial intelligence and machine learning has brought new sets of challenges for business and IT organization's data strategy when it comes to implementing high performance, scalable, and agile cloud-scale architecture.

The next generation of distributed systems for big data analytics need to address data silos between different tiers, such as data lakes, data warehouses, and AI/compute and object storage. It is imperative to develop an infrastructure that sustains healthy data pipelines between storage devices and computing devices (CPUs, GPUs, FPGAs, etc.), reduce network bandwidth, and achieve overall low latency for parallel data processing, all of which is critical for supporting an organization's data-driven goals.

Hadoop ecosystem has evolved over the years from batch processing (Hadoop 1.0) to streaming and near real-time analytics (Hadoop 2.0)

and to Hadoop meets AI (Hadoop 3.0). Currently, the technologies have evolved to enable the data lake as a private cloud, with separation of storage and compute and, going forward, to support hybrid cloud and multi-cloud.

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

- Cloudera Data Platform Private Cloud Base (CDP PVC Base), provides storage and supports 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), provides different experiences or personas (data analyst, data scientist, data engineer) based on processing of workloads for data stored in CDP Private Cloud Base

Apache Ozone provides the foundation for the next generation of storage architecture for Hadoop Distributed File System (HDFS), 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 a specifically designed private cloud for data lake requirements, supporting data intensive workloads with CDP Private Cloud Base, and compute-rich (AI/ML) and compute-intensive workloads with CDP Private Cloud Data Services, while also providing storage consolidation with Apache Ozone on Cisco UCS infrastructure fully managed through Cisco Intersight™. Cisco Intersight simplifies management and moves management of servers from the network into the cloud.

CDIP on CDP Private Cloud is based on Cisco UCS® C245 M6 Rack Server for Storage (Apache Ozone and HDFS) and extends the capabilities of the Cisco UCS rack server portfolio with 3rd Gen AMD EPYC CPUs combined with PCIe 4.0 for peripherals and 3200 MHz DDR4 memory will improve application performance and efficiency.

These servers include the following features:

- Do more with less by taking advantage of up to 64 cores per CPU in 3rd Gen AMD EPYC processors and faster memory performance
- Get the fastest I/O speeds, with PCIe 4.0
- Decrease server Operating Expenses (OpEx) for power and cooling, management, and maintenance with latest generation of M6 servers
- Reduce management complexity with Cisco Intersight Infrastructure Service

We move management from the network into the cloud with Cisco Intersight so that you can respond at the speed and scale of your business and manage all your infrastructure.

CDIP with Cloudera Data Platform enables the customer to independently scale compute and storage resources as needed while offering an exabyte-scale architecture with low Total Cost of Ownership (TCO) and a future-proof architecture with the latest technology offered by Cloudera.

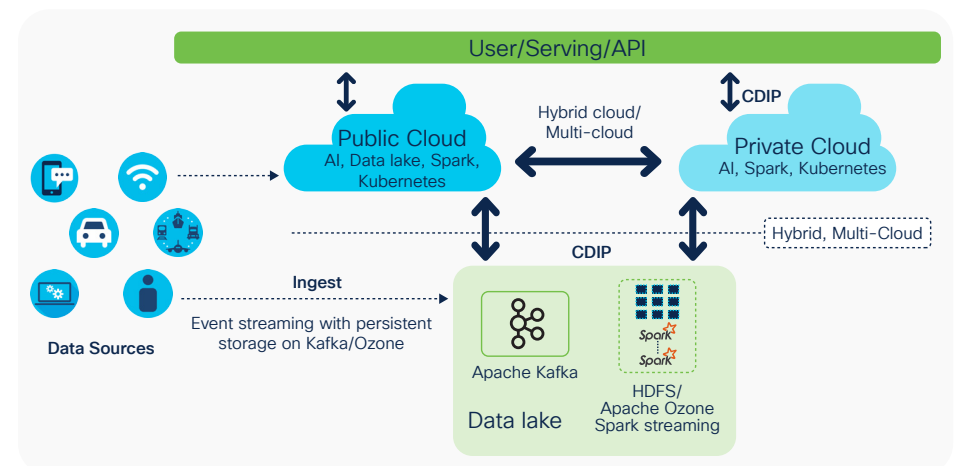
Cisco Data Intelligence Platform

Cisco Data Intelligence Platform (CDIP) is a cloud-scale architecture and a private cloud primarily for a data lake that 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, and FPGA) to work on this data for further analytics
- A storage tier allowing you to gradually retire data that has been worked on to a storage-dense system with a lower \$/TB, providing a better TCO.
- Seamlessly scale the architecture to thousands of nodes with a single pane of glass using Cisco Intersight and Cisco® Application Centric Infrastructure (Cisco 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 Intelligence Platform – journey to hybrid cloud



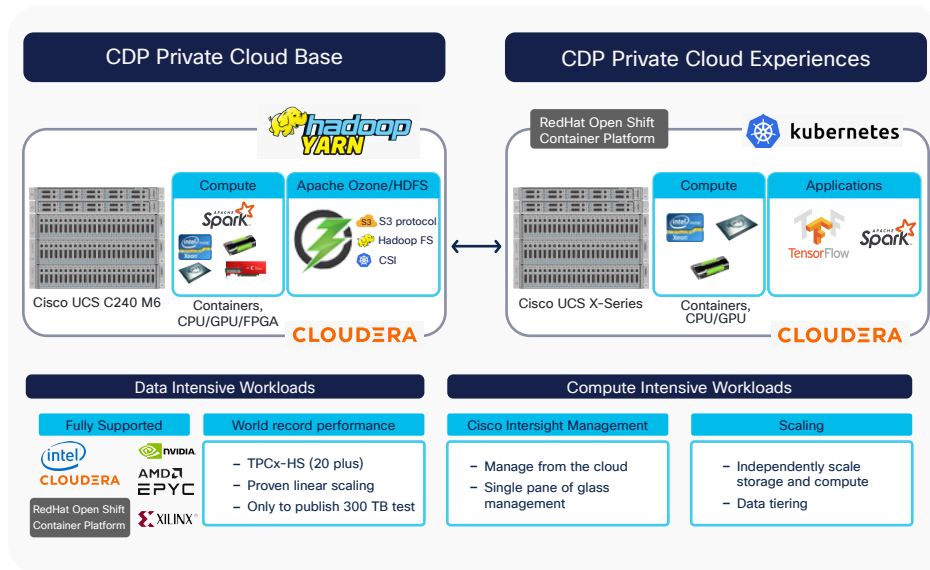
Cisco Data Intelligence Platform with Cloudera Data Platform

Cisco developed numerous industry-leading Cisco Validated Designs (reference architectures) in the area of big data, compute farm with Kubernetes (CVD with Red Hat OpenShift Container Platform) and object store.

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

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

Figure 2. Cisco Data Intelligence Platform with Cloudera Data Platform on Cisco UCS M6



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 Private Cloud (CDP PvC) 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.

Cloudera Data Platform provides:

- **Unified distribution:** Whether you are coming from CDH or HDP, CDP caters to both. It offers richer feature sets and bug fixes with concentrated development and higher velocity.
- **Hybrid and on-premises:** Hybrid and multi-cloud experience, on-premises - it offers the 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:** Policy stick with data, even if the data moves across the entire supported infrastructure.

CDP 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 composed of a variety of components such as 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 preconfigured packages of services are also available for common workloads.

CDP Private Cloud Data Services (CDP PvC DS)

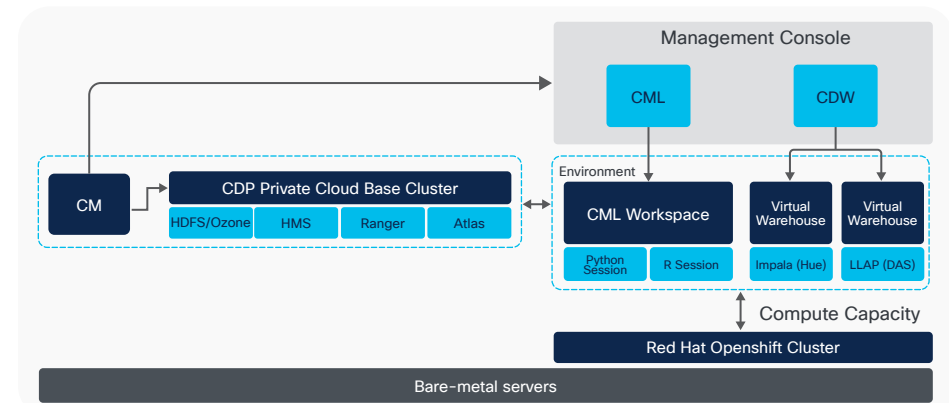
CDP Private Cloud Data Services is the newest on-premises offering of CDP that brings many of the benefits of the public cloud deployments to the on-premises CDP deployments.

CDP Private Cloud provides a disaggregation of compute and storage and allows independent scaling of compute and storage clusters. Using 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 Warehouse and Cloudera Machine Learning services through the management console, and easily scale them up or down as required.

A CDP Private Cloud deployment requires you to have a CDP Private Cloud Base cluster and a Red Hat OpenShift Kubernetes cluster. The OpenShift cluster is set up on a bare-metal deployment. The private cloud deployment process involves configuring a 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 3. CDP Private Cloud Base and CDP Private Cloud Data Services



Cloudera Machine Learning

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

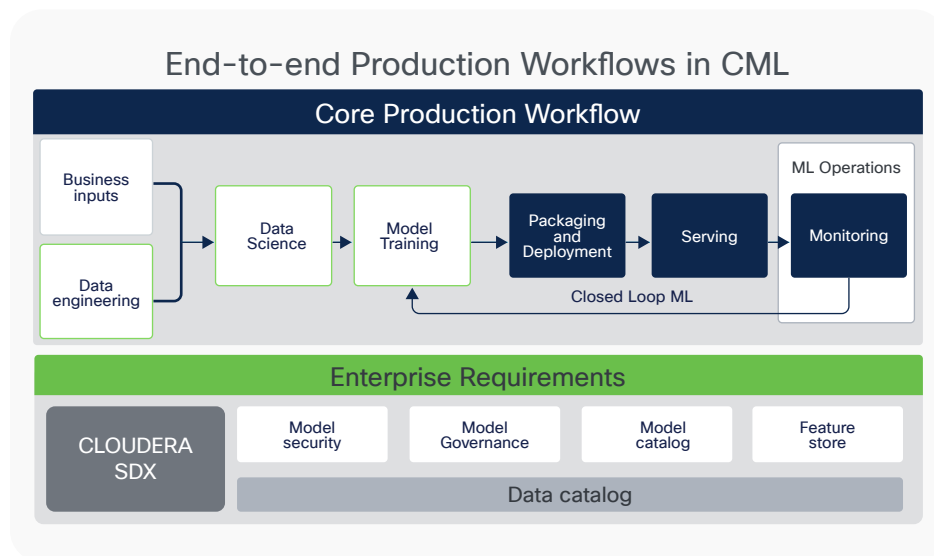
Cloudera Machine Learning (CML) is Cloudera's new cloud-native machine learning service, built for CDP. The CML service provisions clusters, also known as ML workspaces, that run natively on Kubernetes.

Each ML workspace enables 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 enables you to:

- Easily onboard a new tenant and provision an ML workspace in a shared OpenShift environment
- Enable data scientists to access shared data on CDP Private Cloud Base and Cloudera Data Warehouse
- Leverage Spark on Kubernetes to spin up and down Spark clusters on demand

Figure 4. Cloudera Machine Learning (CML)



Apache Ozone

Apache Ozone is a scalable, redundant, and distributed object store for Hadoop. 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 such as Apache Spark, Yarn, and Hive to work natively without any modifications. Apache Ozone is built on a highly available, replicated block storage layer called Hadoop Distributed Data Store (HDDS).


Apache Ozone separates management of namespaces and storage, helping it to scale effectively. Ozone Manager manages the namespaces while Storage Container Manager handles the containers.

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 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.

Figure 5. Data-lake consolidation with Apache Ozone



Solution highlights

- **Intelligent multidomain management with Cisco Intersight** - Enabling IT to operationalize at scale heterogeneous infrastructure and application platform to seamlessly function as a single cohesive unit through a single plane of glass management
- **Powered by latest generation in CPUs from AMD** - Latest generation of processors from AMD (3rd Gen AMD EPYC CPUs) provides the foundation for powerful data center platforms with an evolutionary leap in agility and scalability
- **Eliminate infrastructure silos with CDIP** - A highly modular platform that brings big data, AI compute farms, and object storage to work together as a single entity, while each component can scale independently to address the IT issues in the modern data center
- **Disaggregated architecture** - CDIP is a disaggregated architecture that brings together a more integrated and scalable solution for big data analytics and AI. It is specifically designed to improve resource utilization, elasticity, heterogeneity, and failure handling and also be able to consume

continuously evolving AI/ML frameworks and landscape.

- **Pre-validated and fully supported** - Cisco Validated Designs facilitate faster, more reliable, and more predictable customer deployments by providing configuration and integration of all components into a fully working optimized design, scalability, and performance recommendations.

Fully supported and pre-validated architectural innovations with partners

Pre-tested and pre-validated through industry-standard benchmarks, tighter integration, and performance optimization with industry-leading Independent Software Vendor (ISV) partners in each of these areas: big data, AI, and object storage, Cisco Data Intelligence Platform offers best-of-class end-to-end validated architectures that reduce integration and deployment risk by eliminating guesswork.

For more information, see: https://www.cisco.com/go/bigdata_design

Managing from the cloud with Cisco Intersight

With Cisco Intersight, 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.

Cisco Intersight 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 solution that provides 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 Technical Assistance Center (TAC) enables more efficient support. Cisco Intersight automates uploading files to speed troubleshooting. The Intersight recommendation engine provides actionable intelligence for IT operations management with insights 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. For more details, see: <https://www.cisco.com/c/en/us/products/cloud-systems-management/intersight/index.html>.

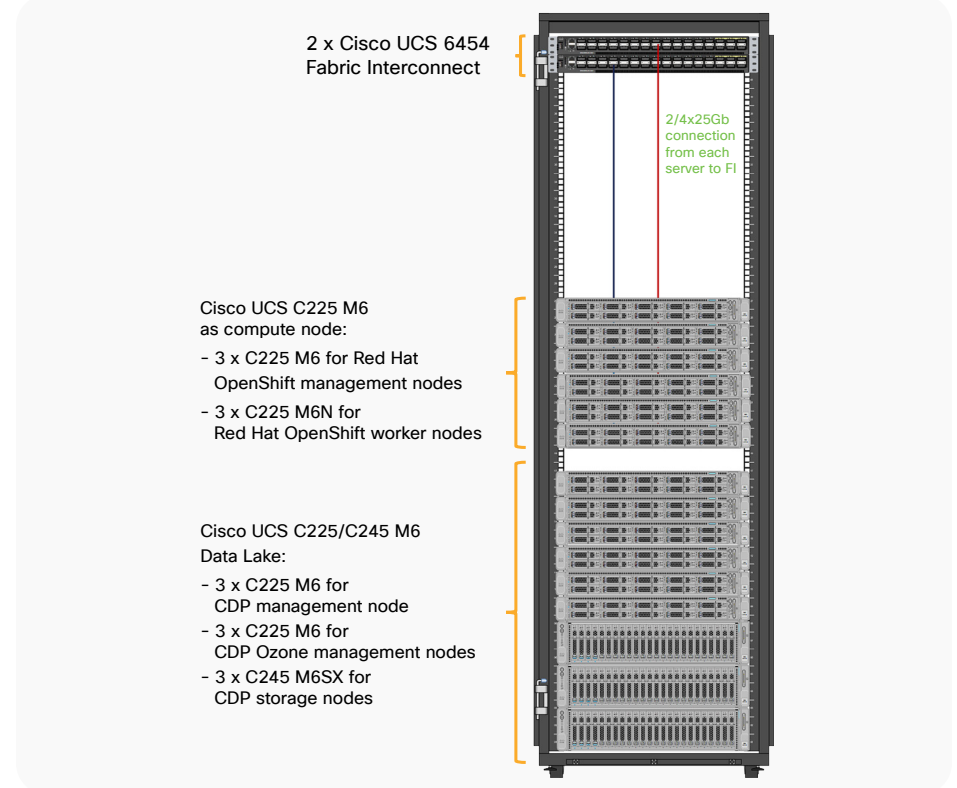
Figure 6 Cisco Intersight



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.

Figure 7. Cisco Data Intelligence Platform with Cloudera Data Platform Private Cloud Data Services



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 rack servers

Component	Performance	High performance
Servers	16 x Cisco UCS C245 M6 Rack Servers with small-form-factor (SFF) drives (UCSC-C245-M6SX)	16 x Cisco UCS C225 M6N Rack Servers with small-form-factor (SFF) drives (UCSC-C225-M6N)
CPU	2 x 3 rd Gen AMD EPYC 7413 (2 x 24 cores @ 2.65GHz)	2 x 3 rd Gen AMD EPYC 7413 (2 x 24 cores @ 2.65GHz)
Memory	16 x 32GB RDIMM DRx4 3200 (512 GB)	16 x 32GB RDIMM DRx4 3200 (512 GB)
Boot	M.2 HWRAID with 2 x 960GB SSDs	M.2 HWRAID with 2 x 960GB SSDs
Storage	24 x 2.4TB 10K rpm SFF SAS HDDs or 24 x 3.8TB 2.5in Enterprise Value 6G SATA SSD and 2x3.8TB 2.5in U.2 NVMe	8 x 7.6TB 2.5in U.2 NVMe and 2x3.8TB 2.5in U.2 NVMe
Virtual Interface Card (VIC)	4x25G mLOM Cisco UCS VIC 1467	4x25G mLOM Cisco UCS VIC 1467
Storage controller	Cisco 12G SAS RAID Controller w/4GB FBWC or Cisco 12-Gbps modular SAS host bus adapter (HBA)	N/A
Network connectivity	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

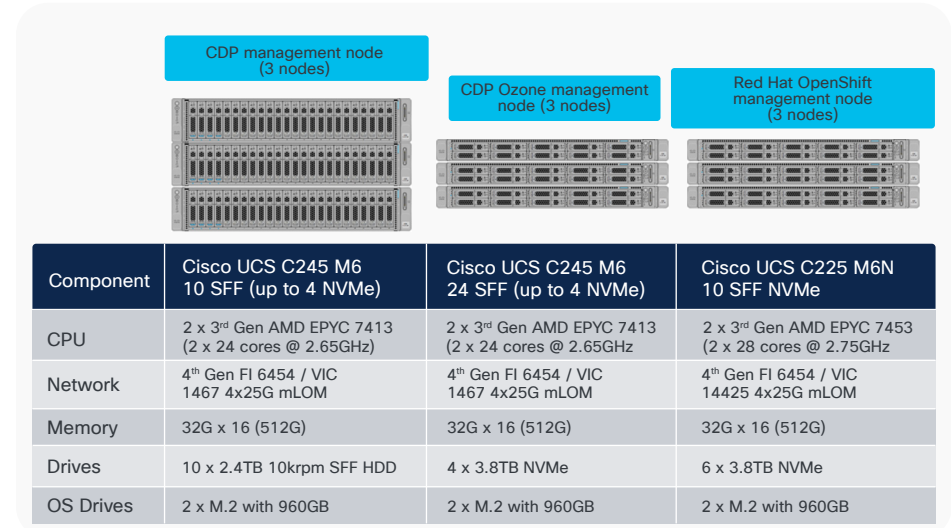
Private cloud reference architecture

Table 2 lists the CDIP private cloud configuration for master and worker nodes.

Table 2. Cisco Data Intelligence Platform with CDP Private Cloud Data Services configuration on Cisco UCS C225 M6 SFF Rack Server

Component	High core option
Servers	Cisco UCS C225 M6SN
CPU	2 x 3 rd Gen AMD EPYC 7453 (2 x 28cores @ 2.75GHz)
Memory	16 x 64GB RDIMM DRx4 3200 (1TB)
Boot	M.2 with 2 x 960GB SSD
Storage	6 x 3.8TB NVMe (Portworx/Ceph [2 drives], local storage [4 drives])
VIC	4 x 25 Gigabit Ethernet with Cisco UCS VIC 14425 mLOM
Network connectivity	Cisco UCS 6454 Fabric Interconnect

Figure 8. Reference architecture for management nodes



Making advanced analytics deployment with CDIP future ready

As enterprises are embarking on the journey of digital transformation, an integrated extensible infrastructure implementation purpose-built to keep pace with constant challenges of technological advancement for each workload can reduce bottlenecks, improve performance, decrease bandwidth constraints, and minimize business disruption.

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack Servers keep pace with 3rd Gen AMD EPYC CPUs innovation by offering the latest processors with higher processor frequency and improved security and availability features. With the increased performance, 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 C225 M6 Rack Server

The Cisco UCS C225 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, EDA, SDS, big data, and edge-centric workloads. You can deploy the Cisco UCS C-Series Rack Servers as standalone servers or as part of the Cisco Unified Computing System™ with the Cisco Intersight Infrastructure Service cloud-based management platform.

The Cisco UCS C225 M6 Rack Server has 3rd Gen AMD EPYC CPUs for the most cores per socket. Combined with PCIe 4.0 for peripherals and 3200 MHz DDR4 memory, you have significant performance and efficiency gains that will improve your application performance.

The C225 M6 Rack Server is single-socket optimized. All I/O is tied to one CPU and its 128 PCIe lanes. Since each server supports up to 2 TB and 64 cores per socket, many customers find that one CPU server now meets their needs. This can reduce software licensing and support costs leading to a better TCO.

Figure 9. Cisco UCS C225 M6 1RU Rack Server



The Cisco UCS C225 M6 Rack Server is designed to deliver exceptional performance, expandability, and efficiency. It offers the following:

- One or two 3rd Gen AMD EPYC CPUs, with up to 64 cores per socket
- Memory:
 - 32 DIMM slots (16 DIMMs per CPU socket), 3200 MHz DDR4 with Up to 4 TB (for 32 x 128 GB DDR4 DIMMs) of capacity
- Up to 10 Small-Form-Factor (SFF) front-loading hot-pluggable drives
 - NVMe/SAS/SATA
- Up to three PCIe 4.0 slots
- Support for VIC 1400 Series and OCP 3.0 network cards
- RAID controller and GPU options are available.
- Internal dual M.2 drive options

Cisco UCS C245 M6 Rack Server

The Cisco UCS C245 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.

You can deploy the Cisco UCS C-Series rack servers as standalone servers or as part of the Cisco Unified Computing System with the Cisco Intersight Infrastructure Service cloud-based management platform. These computing innovations help reduce customers’ Total Cost of Ownership (TCO) and increase their business agility.

These improvements deliver significant performance and efficiency gains that will improve your application performance. The Cisco UCS C245 M6 Rack Server delivers outstanding levels of expandability and performance.

Figure 10. Cisco UCS C245 M6 2RU Rack Server



For more information

- To learn more about Cisco UCS big data solutions, visit: <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/big-data/index.html>
- To find out more about Cisco UCS big data validated designs, visit: https://www.cisco.com/go/bigdata_design
- To learn more about Cisco Data Intelligence Platform, visit: <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 ACI solutions, visit: <http://www.cisco.com/go/aci>

The Cisco UCS C245 M6 Rack Server is designed to deliver exceptional performance, expandability, and efficiency. It offers:

- One or two 3rd Gen AMD EPYC CPUs, with up to 64 cores per socket
- Memory:
 - 32 DIMM slots (16 DIMMs per CPU socket), 3200 MHZ DDR4 with Up to 4 TB (for 32 x 128 GB DDR4 DIMMs) of capacity
- Up to 28 Small-Form-Factor (SFF) front-loading hot-pluggable drives – NVMe/SAS/SATA
- Up to eight PCIe 4.0 slots
- Support for VIC 1400 Series and OCP 3.0 network cards
- RAID controller and GPU options available
- Internal dual M.2 drive options

Conclusion

Evolving workloads need a highly flexible platform to cater to various requirements, whether data-intensive (data lake) or compute-intensive (AI/ML/DL) or just storage-dense (Object Store). In an infrastructure to enable an evolving architecture that could scale to thousands of nodes, operational efficiency can't be an afterthought.

To have a seamless operation of applications at this scale, you need

- Infrastructure automation with centralized management
- Deep telemetry and simplified, granular trouble-shooting capabilities
- Multi-tenancy allowing application workloads that include container and microservices, with the right level of security and SLA for each workload

Cisco UCS with Intersight and Cisco ACI can enable this next-generation cloud-scale architecture deployed and managed with ease.