

Cisco Data Intelligence Platform

Modernizing your data lake in the evolving landscape

In today's digital world, decision making driven by data is quickly becoming the norm. With the accessibility of various toolsets for data collection and analysis, businesses of all sizes can take advantage and drive innovation or define new areas of products or services. Data can be overwhelming when you reach "big data" proportions, making it difficult to glean meaningful insights from it.

Data science, artificial intelligence, and machine learning can lead to the identification of trends, the evaluation of existing programs and the understanding of customer behavior. Leveraging the large and complex data sets ultimately allows adopters to lower costs through better predictions and automation. Implementation of supervised or unsupervised Artificial Intelligence (AI) and Machine Learning (ML) to automate complex problems and analytical tasks from research to production has brought a new set of challenges for enterprises and IT organizations.

The next generation of distributed systems for big data and analytics needs to address data silos between different tiers such as data lakes, data warehouses, AI/compute, and object storage. It is imperative to develop and deploy an end-to-end platform to ingest data, sustain healthy and iterative data pipelines between storage devices and computing devices (CPU, GPU, and FPGA), reduce network bandwidth, and achieve overall low latency for parallel data processing and maintain ML models spanning across multiple use cases.

Cisco® Data Intelligence Platform (CDIP) is a thoughtfully designed private cloud for data lake requirements, supporting data-intensive workloads with Cloudera Data Platform Private Cloud Base, and compute-intensive (AI/ML) workloads with Cloudera Data Platform 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.

Cisco Data Intelligence Platform is based on the Cisco UCS M7 family of servers, which support 4th Gen Intel® Xeon® Scalable Processors with PCIe Gen 5 capabilities supporting 50 percent more cores per socket and 50 percent more memory bandwidth when compared with the previous generation.

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

Cisco Data Intelligence Platform 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) and future-ready architecture supporting complete data life-cycle management.

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
- GPU-accelerated artificial intelligence and data science 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 the gradual retiring of data that has been worked on to a storage-dense system with a lower \$/TB, thus providing a better TCO.
- The capability to seamlessly scale the architecture to thousands of nodes with single pane of glass management using Cisco Intersight and Cisco Application Centric Infrastructure (Cisco ACI®).

The Cisco Data Intelligence Platform caters to an evolving IT architecture for big data and analytics. CDIP brings together a fully scalable infrastructure with centralized management and a 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.

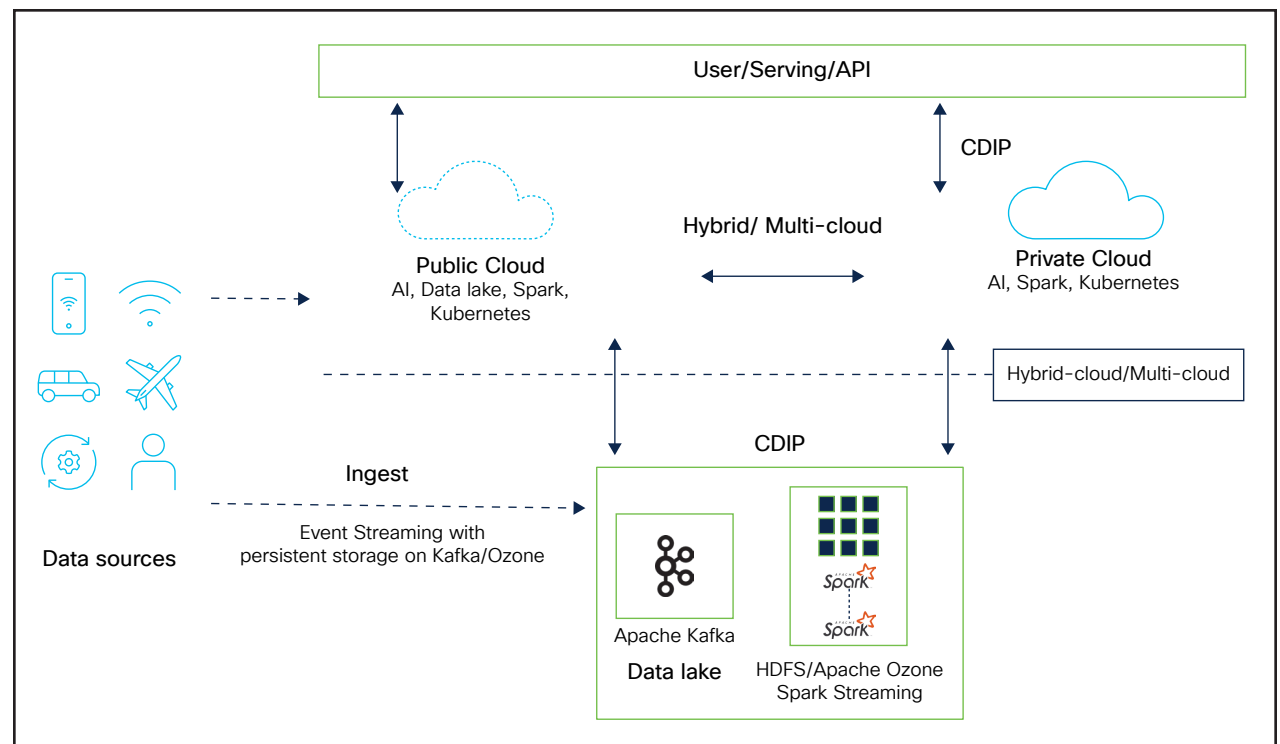


Figure 1. Cisco Data Intelligent Platform with Cloudera Data Platform – data lake evolution to hybrid cloud

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 with the following components:

- Data lake enabled through CDP Private Cloud Base.

- Private cloud with compute on Kubernetes, which can be enabled through CDP Private Cloud Data Services.
- Exabyte storage enabled through Apache Ozone.

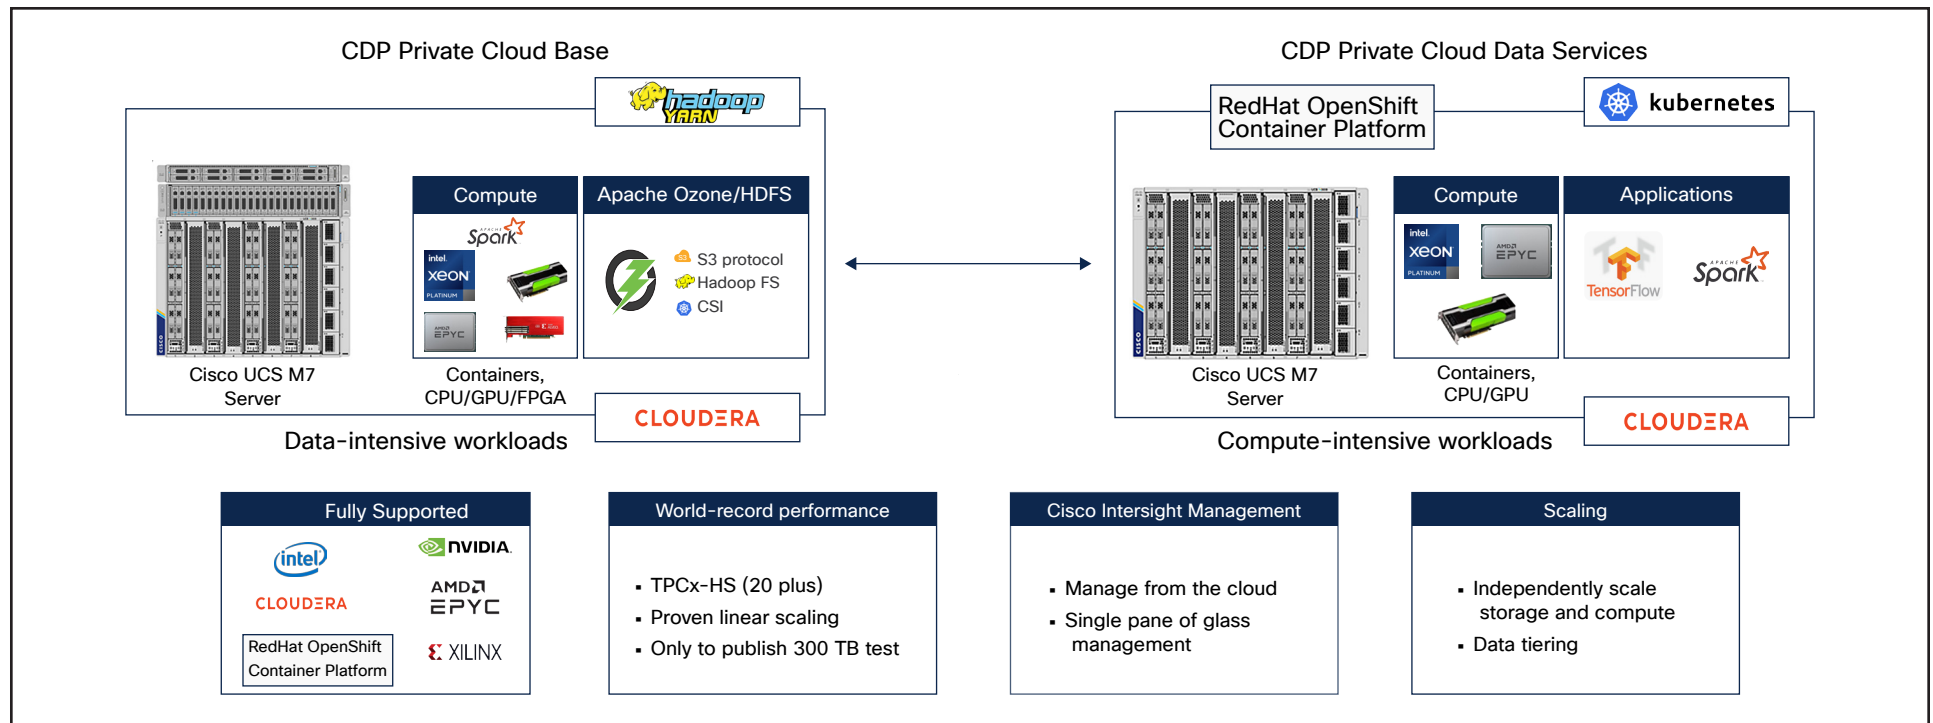


Figure 2. Cisco Data Intelligent Platform with Cloudera Data Platform

Cloudera Data Platform (CDP)

Cloudera Data Platform Private Cloud is the on-premises version of Cloudera Data Platform that is easy to deploy, manage, and use. 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. 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 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 flexibility and on-premises:** on-premises CDP offers industry-leading performance, cost-optimization, and security. It is designed for data centers with optimal

infrastructure. It is Hybrid flexible to deliver use cases in hybrid of on-prem and multi-cloud environments.

- **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:** it provides policy stickiness with data movement across all supported infrastructure.

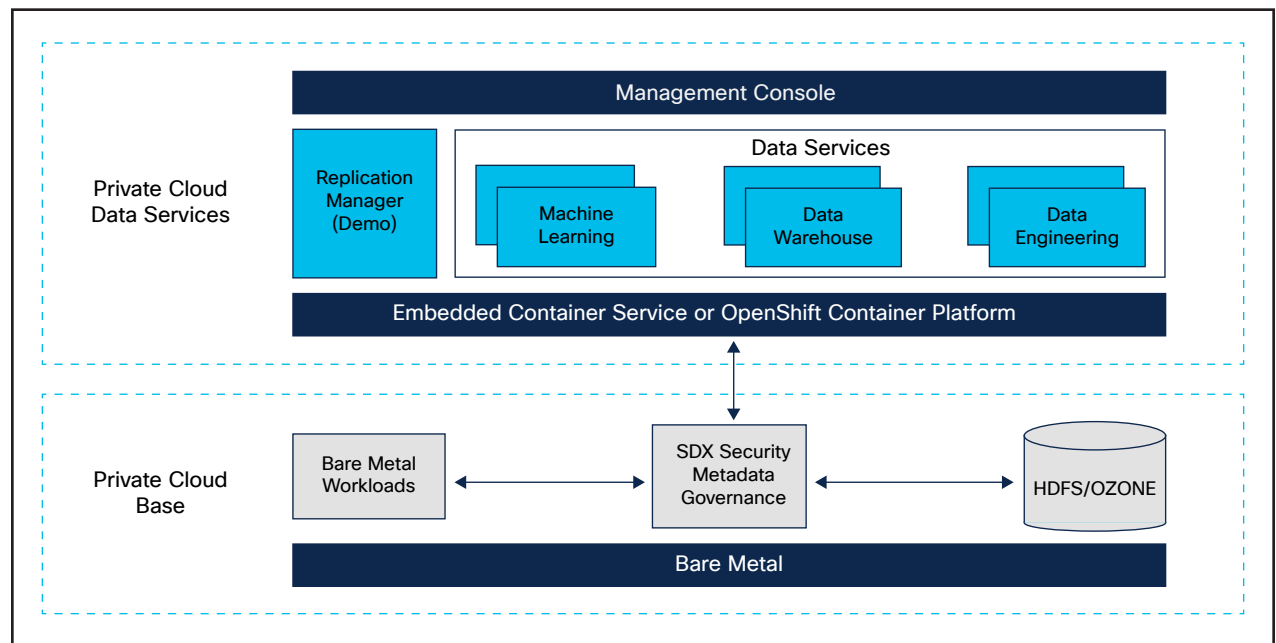


Figure 3. Cloudera Data Platform Private Cloud – high level overview

Cloudera Data Platform 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.

Cloudera Data Platform Private Cloud Data Services deployment requires a Cloudera Data Platform Private Cloud Base cluster and a Kubernetes cluster running on either Red Hat OpenShift or Cloudera's Embedded Container Service (ECS). Cloudera Private Cloud Data Services deployment process involves configuring Cloudera Management Console, registering an environment by providing details of the data lake configured on the Private Cloud Base cluster, and then creating the workloads.

[Cloudera Data Platform Private Cloud Data Services](#) deployment comprises components such as an environment, a [data lake](#), the [Management Console](#) and data services such as [Data Warehouse](#), [Machine Learning](#), [Data Engineering](#), and [Replication Manager](#).

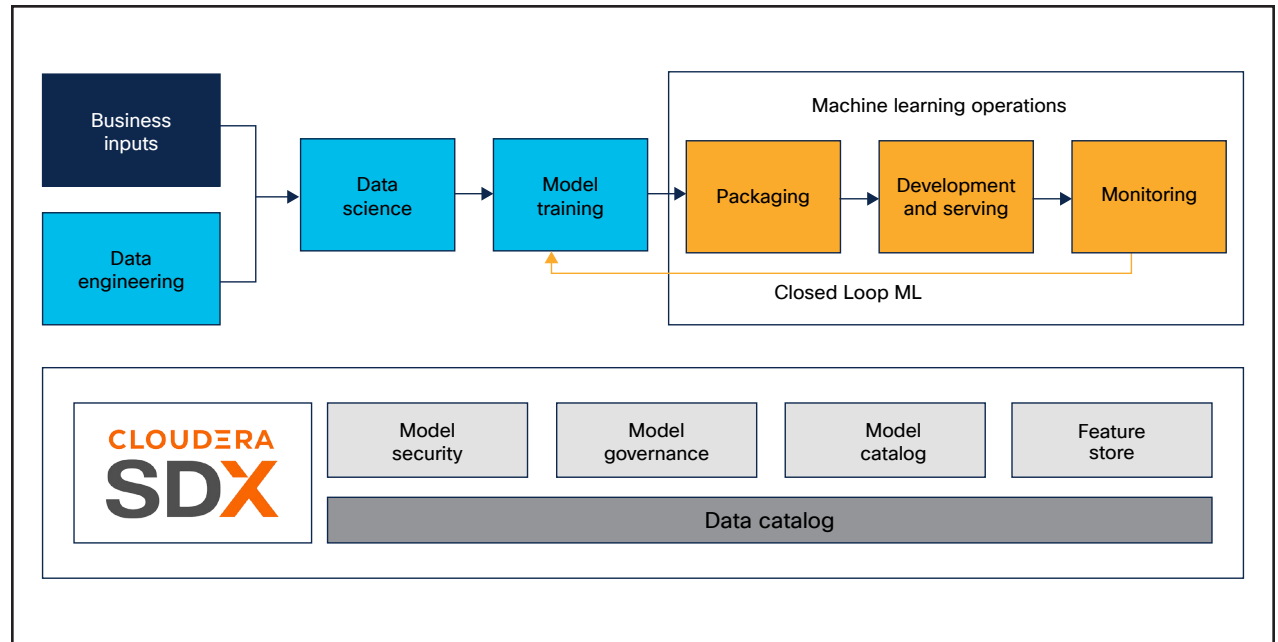


Figure 4. End-to-end production workflow overview in Cloudera Machine Learning

Apache Ozone

Apache Ozone is a scalable, redundant, and distributed object store optimized for big data workloads. 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 work natively without any modifications. Ozone natively supports the S3 API and provides a Hadoop-compatible file system interface. Ozone is typically available in a CDP Private Cloud Base deployment.

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. 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.

Highlights

Intelligent multidomain management with Cisco Intersight

This enables IT to operationalize at scale a heterogeneous infrastructure and application platform to seamlessly function as a single cohesive unit through single pane of glass management.

Powered by latest generation in CPU from Intel

The latest generation of processors from Intel (4th gen Xeon scalable family processors)

provides the foundation for powerful data-center platforms with an evolutionary leap in agility and scalability.

Eliminate infrastructure silos with CDIP

This highly modular platform brings big data, AI compute farms, and object storage to work together as a single entity while each component can scale independently to address 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 and analytics and AI/ML. 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.




Figure 5. Data lake consolidation with Apache Ozone

Fully supported and pre-validated architectural innovations with partners

Cisco Data Intelligence Platform is 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. CDIP offers best-of-class end-to-end validated architectures and reduces integration and deployment risk by eliminating guesswork.

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

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 6. Cisco Data Intelligent Platform with Cloudera Data Platform – reference architecture

CDIP with CDP Private Cloud Base 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 M7 servers

	High performance	Performance	High performance
Servers	Cisco UCS C240 M7SN Rack Servers with small-form-factor (SFF) drives	Cisco UCS C240 M7SX Rack Servers with small-form-factor (SFF) drives	Cisco UCS X210c M7 Compute Node
CPU	2 x 4th Gen Intel Xeon Scalable Processors 6448H processors (2 x 32 cores, at 2.1 GHz)	2 x 4th Gen Intel Xeon Scalable Processors 6448H processors (2 x 32 cores, at 2.1 GHz)	Processors 6448H processors (2 x 32 cores, at 2.1 GHz)
Memory	16 x 32GB 4800 MT/s DDR5 (512GB)	16 x 32GB 4800 MT/s DDR5 (512GB)	16 x 32GB 4800 MT/s DDR5 (512GB)
Boot	M.2 RAID controller with 2 x 960GB SATA SSD	M.2 RAID controller with 2 x 960GB SATA SSD	M.2 RAID controller with 2 x 960GB SATA SSD
Storage	24 x 3.8TB 2.5in U.2 P5520 NVMe High Perf Medium Endurance and 2 x 3.8TB NVMe for Ozone metadata	24 x 2.4TB 10K rpm SFF SAS HDDs or 24 x 3.8TB SATA SSD Enterprise Value and 2 x 3.8TB NVMe for Ozone metadata	6 x 15.3TB 2.5in U.2 P5520 NVMe High Perf Medium Endurance
Virtual Interface Card (VIC)	2 x 100 Gigabit Ethernet with Cisco UCS VIC 15238 mLOM	2 x 100 Gigabit Ethernet with Cisco UCS VIC 15238 mLOM	2 x 100 Gigabit Ethernet with Cisco UCS VIC 15231 mLOM

	High performance	Performance	High performance
Storage controller	<ul style="list-style-type: none"> 2 x Cisco Tri-Mode 24G SAS RAID Controller w/4GB Cache* 	2 x Cisco Tri-Mode 24G SAS RAID Controller w/4GB Cache* or Cisco 12G SAS RAID controller w/4GB FBWC	Passthrough controller
Network connectivity	Cisco UCS 6536 Fabric Interconnect	Cisco UCS 6536 Fabric Interconnect	Cisco UCS 6536 Fabric Interconnect
GPU (optional)	NVIDIA GPU A16 or A30 or A100-80	NVIDIA GPU A16 or A30 or A100-80	NVIDIA GPU T4, A16 or A100-80

* Tri-mode controller will be available post first customer ship (FCS).

CDIP with CDP Private Cloud Data Services reference architecture

Table 2 lists the CDIP with CDP Private Cloud Data Services cluster configuration (Red Hat OpenShift Container Platform (RHOCP) or Cloudera Embedded Container Service (ECS)).

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

	High-core option
Servers	Cisco UCS X210c M7 Compute Node (up to 8 per chassis)
CPU	2 x 4th Gen Intel Xeon Scalable Processors 6448H processors (2 x 32 cores, at 2.1 GHz)
Memory	16 x 64GB 4800 MT/s DDR5 (1024Gb)
Boot	M.2 RAID controller with 2 x 960GB SATA SSD

	High-core option
Storage	6 x 3.8TB 2.5in U.2 P5520 NVMe High Perf Medium Endurance
Virtual Interface Card (VIC)	2 x 100 Gigabit Ethernet with Cisco UCS VIC 15231 mLOM
GPU (optional)	Cisco UCS X440p PCIe node with UCS PCI Mezz Card for X-Fabric and NVIDIA GPU T4, A16 or A100-80
Network connectivity	Cisco UCS 6536 Fabric Interconnect

Future proofing advanced analytics deployment with CDIP

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.

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). With an infrastructure to enable

this evolving architecture that can scale to thousands of nodes, operational efficiency can't be an after-thought.

To bring in seamless operation of applications at this scale, one needs:

- Infrastructure automation with centralized management.
- Deep telemetry and simplified granular trouble-shooting capabilities.
- Multitenancy, allowing application workloads including containers 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.

For more information

For additional information, see the following resources:

- [To find out more about Cisco UCS big data solutions.](#)
- [To find out more about Cisco UCS big data validated designs.](#)
- [To find out more about Cisco UCS AI/ML solutions.](#)
- [To find out more about Cisco Data Intelligence Platform.](#)