

Data Management with Cloudera Data Platform on Intel-powered Dell EMC Infrastructure

Design Guide

Abstract

This design guide provides architecture and design information for the Dell Technologies Validated Design for Data Management with Cloudera Data Platform (CDP) Private Cloud Base, for deployment on Intel-powered Dell EMC PowerEdge servers, PowerSwitch networking, and PowerScale storage.

Dell Technologies Solutions

Dell Technologies

Validated Design

Copyright

© 2020 - 2021 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

Revision history.....	5
Chapter 1: Introduction.....	6
Solution introduction.....	6
Audience.....	6
What's new.....	7
Upgrade enhancements.....	7
General feature enhancements.....	7
Design Guide introduction.....	7
Chapter 2: Architecture concepts and requirements.....	9
Introduction.....	9
Node-level architecture.....	9
Role assignment recommendations.....	10
Node-level architecture with PowerScale storage	11
High availability.....	12
Cluster sizing and scaling.....	12
Cluster scaling.....	13
Cluster storage sizing.....	14
Worker node sizing.....	15
Master node sizing.....	15
PowerScale sizing and scaling.....	15
Solution requirements.....	15
Install the GPU driver.....	16
Chapter 3: Network architecture.....	17
Introduction.....	17
Cluster networks.....	18
Server node connections.....	18
Network high availability.....	19
Network fabric.....	19
Pod switches.....	19
Aggregation switches.....	21
Layer 3 cluster aggregation.....	21
iDRAC management network.....	22
Core network integration.....	22
PowerScale network integration.....	22
25 GbE network equipment summary.....	23
Chapter 4: Infrastructure configurations.....	25
Infrastructure nodes.....	25
Worker nodes.....	27
General purpose Worker nodes.....	27
All flash Worker nodes.....	29
GPU accelerated Worker nodes.....	31

PowerScale Worker nodes	33
PowerScale node.....	34
Chapter 5: Software infrastructure.....	36
Software components.....	36
Chapter 6: Validation.....	37
Overview.....	37
Component validation.....	37
TeraSuite test.....	37
DSFIO test.....	38
MapReduce test.....	38
Spark test.....	39
Spark GPU test.....	39
Hive test.....	40
HBase test.....	40
Findings.....	41
Summary.....	41
Chapter 7: Summary.....	42
Overview.....	42
We value your feedback.....	42
Chapter 8: References.....	43
Dell Technologies documentation.....	43
Cloudera documentation.....	43
Services engagements.....	44
Dell Technologies Customer Solution Centers.....	44
Dell EMC OpenManage Enterprise Power Manager.....	45
Dell Technologies Info Hub.....	45
More information.....	45

Revision history

Table 1. Document revision history

Date	Document revision	Description of changes
October 2021	1.2	Updated: <ul style="list-style-type: none">● Cloudera CDP Private Cloud Base software version from 7.1.5 to 7.1.7● Introduction● Architecture concepts● Network infrastructure● Server infrastructure● Software and firmware versions● References Added: <ul style="list-style-type: none">● Validation
March 2021	1.1	Updated: <ul style="list-style-type: none">● Name change from CDP Data Center to CDP Private Cloud Base● Cluster pricing model● Hardware architecture
July 2020	1.0	Initial release

Introduction

Cloudera Data Platform (CDP) Private Cloud Base, previously known as CDP Data Center, is the on-premises version of Cloudera Data Platform. This product combines the best of the prior Cloudera's and Hortonworks' technologies, along with new features and enhancements. CDP Private Cloud Base is also used with CDP Private Cloud Data Services to form CDP Private Cloud.

Topics:

- [Solution introduction](#)
- [What's new](#)
- [Design Guide introduction](#)

Solution introduction

CDP Private Cloud delivers powerful analytic, transactional, and machine learning workloads in a hybrid data platform. It combines the agility and flexibility of a public cloud with the control of the data center. With elastic analytics and scalable object storage, CDP Private Cloud modernizes traditional single-cluster deployments into a scalable and efficient end-to-end data platform.

CDP Private Cloud Base is the unification of Cloudera Distribution for Apache Hadoop (CDH) and Hortonworks Data Platform (HDP), giving customers the best of both worlds. This new product combines the best technologies from Cloudera and Hortonworks with new features and enhancements across the stack. CDP Private Cloud Base forms a comprehensive data platform that encompasses the entire data life cycle. This unified distribution is a scalable and customizable platform where you can securely run many types of data analytics workloads.

CDP Private Cloud Base can be a stand-alone data analytics platform. It also supports a hybrid or multicloud solution, where compute tasks can be separated from data storage, and where data can be accessed from remote clusters. This approach provides a foundation for containerized applications by managing storage, table schema, authentication, authorization, and governance in CDP Private Cloud Base. It consists of various 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.

In summary, CDP Private Cloud Base is a stand-alone instance of Cloudera Data Platform for the on-premises data center that can also be deployed with the CDP Private Cloud Data Services cluster to form the complete CDP Private Cloud. CDH and HDP customers are encouraged to upgrade to CDP Private Cloud Base for improved enterprise data management capabilities and new platform innovations.

Audience

This document is intended for data center managers and IT architects who are involved with designing, planning, or operating the hardware and software infrastructure for CDP Private Cloud, for:

- New deployments
- Upgrades or migrations from:
 - CDH
 - HDP

This document assumes some familiarity with Cloudera Data Platform capabilities and functions.

NOTE: This document may contain language from third-party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

What's new

CDP Private Cloud Base 7.1.7 introduces key new features to improve the management, security, and analytical capabilities of the platform and provides additional platform support. It is a cumulative maintenance release that carries forward the improvements from CDP Private Cloud Base 7.1.6 and prior releases.

In addition to new features and platform support, an important goal of the last few releases has been to make the transition to CDP easier by adding more upgrade options. Below is a brief overview of the new features and improvements that were added in this release.

Upgrade enhancements

To help customers with their move to CDP Private Cloud Base, this release contains the following upgrade enhancements:

- In-place upgrades are now available for CDH versions 6.1.x, 6.2.x, and 6.3.x., completing the in-place upgrade vision from all supported CDH and HDP versions.
- Documented rollback procedures for CDH 6 to CDP Private Cloud Base 7.1.7 and for HDP 3 to CDH 7.1.6.
- An [upgrade companion](#) to complement existing upgrade documentation that guides you through your upgrade journey.

General feature enhancements

To further improve the management, security, and analytical capabilities of the platform, the following updates are in this release:

- Several enhancements have been made to Cloudera Manager including security fixes and the removal of all common vulnerabilities and exposures (CVEs) across over 20 embedded libraries. These enhancements make version 7.1.7 the most secure release of CDP Private Cloud Base yet. Also added are new features and tools for improved platform efficiency and ease of use.
- Many updates have been made to SDX for improved platform and data governance. These enhancements include a more secure method of authorizing Hive table creation, auditing HDFS superuser actions in Ranger, Atlas and Kafka integration, Impala Row Filtering, and many other features.
- Data Warehouse enhancements include support for Ranger-based row-level filter policies in Impala. Row-level filters are similar to other Ranger access policies and can be set for specific users, groups, and conditions.

These updates are some of the highlights of this release and not a comprehensive list of new features. See the [CDP Private Cloud Base 7.1.7 Release Summary](#) for the full list and details of these features.

Design Guide introduction

This document is the Design Guide for the Dell Technologies Validated Design (DTVD) for Data Management with Cloudera CDP Private Cloud Base. It may be read along with the associated white paper of the same name, which provides an overview of what an enterprise data platform is, benefits, and typical use cases. The white paper provides a description of Cloudera Data Platform, including the component clusters of CDP Private Cloud Base and CDP Private Cloud Data Services, and discusses the journey to CDP, including:

- Upgrades and migrations to CDP Private Cloud Base
- The relation of CDP Private Cloud Base as a foundation for CDP Private Cloud

This design guide provides the infrastructure guidance and a validated reference architecture, which is designed for high availability, that includes:

- The software infrastructure and cluster architecture
- The Dell EMC PowerEdge server configurations
- The Dell EMC PowerSwitch networking architecture and configurations

When CDP Private Cloud Base is being used with CDP Private Cloud Data Services to form CDP Private Cloud, this document should be used with the [CDP Private Cloud Data Services documentation](#) on the [Cloudera documentation website](#).

Dell Technologies and Cloudera have been collaborating for over eight years to provide customers with guidance on optimal hardware to streamline the design, planning, and configuration of their Cloudera deployments. This document is based on the collective experience of both companies in deploying and running enterprise production environments.

Architecture concepts and requirements

This chapter describes the cluster architecture and solution requirements that Dell Technologies recommends for Cloudera CDP Private Cloud Base.

Topics:

- [Introduction](#)
- [Node-level architecture](#)
- [Node-level architecture with PowerScale storage](#)
- [High availability](#)
- [Cluster sizing and scaling](#)
- [Solution requirements](#)

Introduction

Cloudera CDP Private Cloud Base provides data management, enterprise analytics, and management tools for big data. The data management services include HDFS file storage and Ozone object storage. The Cloudera Runtime provides the analytics services, which include components like Hive, HBase, MapReduce, and Spark. The management tools include:

- Cloudera Manager for cluster management, monitoring, and configuration
- Cloudera SDX for security, governance, and metadata

Successful deployment and operation of Cloudera CDP Private Cloud Base depends on a well-designed infrastructure architecture. The architecture must provide high performance, scalability, reliability, and manageability.

Node-level architecture

[Node architecture](#) depicts the nodes and services in a CDP Private Cloud Base deployment.




Figure 1. Node architecture

The cluster infrastructure consists of multiple physical server nodes. Each node has a physical configuration that is specialized for its role in the cluster. These nodes are further specialized through the software services that are assigned to them. The table below defines the various cluster nodes and their physical configuration.

Table 2. Node definitions

Node	Definition	Physical configuration
Master node	This node runs all the services that are required to manage the cluster storage and compute services.	Infrastructure node
Utility node	This node runs Cloudera Manager and the Cloudera Management Services.	Infrastructure node
Edge node	This node contains all client-facing configurations and services, including gateway configurations.	Infrastructure node
Worker node	This node runs all the services that are required to store blocks of data on the local hard drives, and run processing tasks against that data.	Worker node

A high-performance network fabric connects the cluster nodes in a Cluster Data network. An additional Edge network provides an interface between the cluster and external systems and applications. See [Network architecture](#) for a detailed description of the network design.

The minimum supported configuration is eight cluster nodes, which include three Master nodes, one Utility node, one Edge node, and three Worker nodes. Dell Technologies recommends a ten-node cluster with five Worker nodes as a starting point.

Role assignment recommendations

The table below describes recommended host role assignments in a medium-sized high availability deployment.

Table 3. CDP Private Cloud Base nodes and roles

Node	Service
Master node 1	NameNode, JournalNode, FailoverController, YARN ResourceManager, ZooKeeper, JobHistory Server, Spark History Server, Kudu master
Master node 2	NameNode, JournalNode, FailoverController, YARN ResourceManager, ZooKeeper, Kudu master
Master node 3	JournalNode, ZooKeeper, Kudu master (All require an odd number of masters for high availability.)
Utility node 1	Cloudera Manager, Cloudera Management Service, Hive Metastore, Impala Catalog Server, Impala StateStore, Oozie, ZooKeeper (requires a dedicated disk), JournalNode (requires a dedicated disk), Apache Atlas, Apache Ranger
Edge nodes	Hue, HiveServer2, Gateway configuration
Worker nodes	DataNode, NodeManager, Impala Daemon (impalad), Kudu tablet server

These recommendations for role assignments are intended as a starting point. The role assignments may differ depending on the cluster size and the services that are used. See [Runtime Cluster Hosts and Role Assignments](#) in the CDP Private Cloud Base documentation for more details.

Node-level architecture with PowerScale storage

An alternative cluster architecture can be used that uses PowerScale OneFS HDFS protocol for storage. The figure below depicts the nodes and services in this deployment scenario.




Figure 2. Node-level architecture with PowerScale storage

Compute and storage can be scaled independently using this alternative architecture. The Isilon storage nodes provide the HDFS NameNode and DataNode services instead of the services being assigned to the Master nodes and Worker nodes. The Worker nodes only include enough storage for runtime operations like shuffle-sort spill files and cache.

This alternative architecture reduces the HDFS bandwidth requirements for the Cluster Data network. PowerScale OneFS implements data durability internally and uses a private back-end network for internal operations. A single copy of the data is transferred to the Isilon storage nodes when Worker nodes write to HDFS. No replication traffic occurs on the Cluster Data network. Also, HDFS recovery traffic for failed drives or nodes does not occur on the Cluster Data network.

High availability

This design implements high availability (HA) at multiple levels through a combination of hardware redundancy and software support.

Hadoop storage resiliency	CDP HDFS implements data resiliency through replication and erasure coding. The HDFS implementation understands node and rack locality, and distributes data to minimize the impact of a node- or rack-level failure.
HDFS highly available NameNodes	This design implements high availability for the HDFS directory through a quorum mechanism that replicates critical NameNode data across multiple physical nodes.
Network resiliency	The production network can optionally use bonded connections to pairs of switches in each pod. Pod-level switches should use redundant connections to switch pairs at the aggregation level. This configuration provides increased bandwidth capacity and allows operation at reduced capacity if a network port, network cable, or switch fails.
Resource manager high availability	This design supports high availability for the Hadoop YARN resource manager. Without resource manager HA, a Hadoop resource manager failure causes running jobs to fail. If there is a resource manager failure, jobs can continue running when resource manager HA is enabled.  NOTE: Dell Technologies recommends resource manager HA for production clusters.
Database server high availability	This design supports high availability for the operational databases. The database server that is used for the Cloudera Manager operational and metadata databases stores its data on a RAID 6 or 10 partition. If there is a drive failure, that partition provides storage reliability. Database server high availability can be implemented using: <ul style="list-style-type: none">• One or more additional PostgreSQL instances on other nodes in the cluster• An external, high availability database server  NOTE: Dell Technologies recommends implementing HA for the database server.

Cluster sizing and scaling

Cluster sizing and scaling are two different but related considerations. Sizing is concerned with ensuring the cluster meets the workload requirements for storage and processing throughput. Scaling is concerned with growth of the cluster over time as capacity needs increase. Since Cloudera CDP Private Cloud Base is a parallel scale-out system, some sizing requirements can be addressed through scaling while others must be addressed through node level sizing.

Sizing and scaling of a CDP Private Cloud Base cluster are complex topics that require knowledge of the workloads. This section highlights the main considerations that are involved but does not provide detailed recommendations. Your Dell Technologies or authorized partner sales representative can help with detailed sizing calculations.

There are many parameters that are involved in cluster sizing. The primary parameters are:

Storage capacity	Storage capacity is usually the first parameter that is used to size a cluster. Calculating storage capacity is important and straightforward. However, storage capacity should be calculated while taking the other sizing parameters into account to maintain a balance between storage and processing capability.
Data ingest volumes and growth rates	Data ingest volume and growth rates each have multiple impacts on cluster sizing. Storage capacity should account for growth due to data ingest and growth of ingest volumes. Data ingest also impacts network utilization of the Edge and Cluster Data networks. The network bandwidth that is required for ingest, and the amount of ingest processing required, determine the number of Edge nodes that the cluster requires.

Memory and processor capacity

Memory and processor requirements for jobs running on the cluster must be considered when sizing. Memory and processor capacity increases as nodes are added to the cluster. Workloads like HBase or Spark jobs with large memory requirements may impact sizing of individual nodes in addition to the overall size of the cluster.

Service level agreements

Production cluster sizing must meet any performance requirements that service level agreements (SLAs) specify. If there are critical path jobs that must meet a specific execution time or throughput, the cluster sizing and balance between compute and storage may have to be adjusted accordingly. Overall cluster throughput is as important as storage capacity, and often influences the number of nodes independent of the required storage capacity.

Cluster scaling

The architecture is organized into three units for scaling as the CDP Private Cloud Base environment grows. From smallest to largest, they are:

- Rack
- Pod
- Cluster

Each unit has specific characteristics and sizing considerations that are documented in this guide. The architecture design enables you to scale the CDP Private Cloud Base environment by adding additional capacity as needed, without replacing any existing components.

Rack

A rack is the smallest size designation for a CDP environment. A rack consists of the power, network cabling, and data and management switches to support a group of Worker nodes.

A rack is a physical unit. Physical constraints define its capacity, including available space, power, cooling, and floor loading. A rack should use its own power within the data center, independent from other racks, and should be treated as a fault zone. If a rack level failure occurs in a multiple-rack pod or cluster, the cluster continues to function with reduced capacity.

This architecture uses 12 nodes as the nominal size of a rack, but higher or lower densities are possible. Typically, a rack contains about 12 nodes using scale-out servers like the Dell EMC PowerEdge R650 and Dell EMC PowerEdge R750. The node density of a rack does not affect overall cluster scaling and sizing, but it does affect fault zones in the cluster.

Pod

A pod is the set of nodes that are connected to the first level of network switches in the cluster. It consists of one or more racks. A pod can include a smaller number of nodes initially, and expand over time to a maximum that the available switch ports define.

A pod is a second-level fault zone above the rack level. If a pod level failure occurs in a multiple pod cluster, the cluster continues to function with reduced capacity. A pod can support enough server nodes and network switches for a minimum commercial scale installation.

In this architecture, a pod supports up to 48 nodes (nominally four racks). The recommended pod size is 36 nodes. The size of a pod can vary from this baseline recommendation. Changing the pod size affects the network bandwidth oversubscription at the pod level, the size of the fault zones, and the maximum cluster size.

Cluster

A cluster is a single CDP environment that is attached to a pair of network switches providing an aggregation layer for the entire cluster.

A cluster can range in size from a single pod in a single rack to many pods in multiple racks. A single pod cluster is a special case and can function without an aggregation layer. This scenario is typical for smaller clusters before additional pods are added.

At the cluster level, pod-to-pod (or uplink) network bandwidth factors into overall cluster performance. The maximum oversubscription ratio is of uplink bandwidth 4:1. For most clusters, the oversubscription ratio should be kept close to 2:1. The number of nodes in each pod defines the oversubscription ratio.

Cluster node counts

The minimum supported configuration is eight nodes:

- Three Master nodes
- One Utility node
- One Edge node
- Three Worker nodes

Although a minimum of one Edge node is required per cluster, larger clusters and clusters with high ingest volumes or rates may require additional Edge nodes. Cloudera recommends a baseline of one Edge node for every 20 Worker nodes.

The recommended cluster node counts are 36 nodes per pod, which results in a 2.25:1 oversubscription ratio and a maximum cluster size of 1152 nodes. [Recommended cluster sizing](#) shows the recommended number of nodes per pod and pods per cluster. It also shows some alternatives for cluster sizing with different bandwidth oversubscription ratios.

Table 4. Recommended cluster sizing

Nodes per rack	Nodes per pod	Pods per cluster	Nodes per cluster	Bandwidth oversubscription
12	36	32	1152	2.25:1
12	48	32	1536	3 :1
12	24	32	768	1.5 :1
12	24	32	1024	2 :1

Cluster storage sizing

Total cluster storage capacity is a function of the server platform and disk drives chosen, and scales with the number of Worker nodes.

The amount of usable storage in a cluster also depends on the types of data durability and data compression that are used. The usable storage capacity can be calculated as:

`number of nodes x raw storage per node x storage efficiency x compression ratio`

This calculation is straightforward but depends on estimating the storage efficiency and compression ratio.

The Hadoop Distributed File System (HDFS) storage system supports two options for data durability: replication and erasure coding. These options have different storage efficiencies.

When replication is used, HDFS creates multiple copies of data across nodes to guard against data loss. The number of replicas (replication factor) is configurable and can be changed on a file-by-file basis. The default replication factor is three. That value that is typically used for storage capacity estimates.

HDFS replication decreases the storage efficiency by the replication factor.

When erasure coding is used, data is divided into blocks, encoded with parity, and distributed across nodes. The details of the encoding are specified in an erasure coding policy. Erasure coding policies allow a tradeoff between data durability and storage efficiency. For example, a Reed-Solomon 6-3 policy has durability of 3 and a 67% storage efficiency. A Reed-Solomon 3-2 policy has durability of 2 and a 60% storage efficiency.

HDFS replication and erasure coding can be used simultaneously in a cluster. Erasure coding policies are specified at the HDFS directory level, while replication is specified on a file-by-file basis.

Compression can also be used to reduce the storage required. Compression is optional and applies to individual files. HDFS supports multiple data compression codecs, and it is possible for each compressed file to use a different codec. The compression ratio that is achieved for a given file depends on both the dataset and codec that is used, and is difficult to estimate. The best approach is to:

1. Test several different codecs on real datasets to determine what works best.
2. Make estimates on the amount of data to be compressed.

Worker node sizing

Worker node configurations can be optimized for specific workloads.

The recommended configurations are designed to provide a balance of disk and network throughput for general-purpose workloads.

NOTE: Dell Technologies does not recommend reducing the number of disks in the Worker node. Doing so reduces the available storage bandwidth.

Changes to processor and memory size can be made to handle specific workloads. The impact on the overall cluster capacity should also be considered when changing the configurations. See [Worker nodes](#) for more details.

Master node sizing

The hardware configurations for the Infrastructure nodes support clusters in the petabyte storage range without changes. You can add more Infrastructure nodes to distribute services across multiple nodes for increased capacity. Adding Infrastructure nodes may be necessary for operational databases or to implement dedicated servers for services like ZooKeeper and Ranger. See [Infrastructure nodes](#) for more details.

PowerScale sizing and scaling

If PowerScale is used for the primary HDFS storage, the sizing and scaling of the cluster is different from implementations using servers with direct attached storage.

The primary difference is that scaling of compute and storage is performed independently by adding either PowerEdge Worker nodes or PowerScale storage nodes. This difference simplifies changing the balance of compute and storage over time. Dell Technologies recommends using the Dell EMC PowerScale Isilon H5600. A ratio of five Worker nodes to one Isilon H5600 storage node is a good starting point.

NameNode sizing is also different. When you use PowerScale for primary HDFS storage, the NameNode services are integrated into the PowerScale nodes. Separate NameNode infrastructure servers are not required to be sized and scaled, and the NameNode services are scaled as PowerScale storage nodes are added.

Network bandwidth utilization is different with PowerScale HDFS storage. PowerScale's data durability replaces HDFS replication and erasure coding when using PowerScale. OneFS implements data durability internally and uses a dedicated backend network. A uniform amount of network bandwidth is used for both reads and writes with PowerScale HDFS because HDFS write replication does not consume additional network bandwidth. This option simplifies cluster throughput calculations and increases performance.

Storage sizing calculation is different since PowerScale determines the translation from raw storage capacity to usable HDFS capacity, rather than HDFS replication and erasure coding.

Some local storage is required on the Worker nodes for intermediate storage when using PowerScale HDFS. This storage is used for MapReduce spill files, Spark cache, and other temporary files. The total intermediate storage across the cluster should be approximately 20% of the total PowerScale HDFS storage. The recommended PowerScale Worker node configurations include this storage. See [PowerScale Worker nodes](#) for more information.

Solution requirements

Data Management with Cloudera Data Platform on Intel-powered Dell EMC Infrastructure may have prerequisites for some features. For this version, see:

- [Install the GPU driver](#)

Install the GPU driver

About this task

 **NOTE:** This procedure is only required for GPU-accelerated servers.

The Linux kernel that is available with Red Hat Enterprise Linux Server (RHEL) 8.2 does not include a prebuilt binary driver and CUDA environment from NVIDIA. If you attempt to perform a `dnf install` of the NVIDIA driver, the server updates to the RHEL 8.4 release kernel. However, Cloudera CDP Private Cloud Base supports RHEL 8.2. NVIDIA has not yet released an RPM driver and CUDA combination. You can download the latest source package from the [NVIDIA developer website](#) and build a compatible driver and CUDA environment.

Steps

1. Install the developer files and kernel headers:

```
# dnf install kernel-devel kernel-headers gcc make elfutils-libelf \
elfutils-libelf-devel libglvnd libglvnd.i686
```

2. Download the source package:

```
# wget https://developer.download.nvidia.com/compute/cuda/11.4.2/local_installers/\
cuda_11.4.2_470.57.02_linux.run
```

3. Run the executable to install the driver:

```
# sh cuda_11.4.2_470.57.02_linux.run
```

Results

The RHEL 8.2-compatible CUDA driver is now installed.

Network architecture

Dell EMC PowerSwitch products provide a fast, robust, and proven network architecture that is well suited for Cloudera CDP Private Cloud Base clusters.

Topics:

- [Introduction](#)
- [Cluster networks](#)
- [Network fabric](#)
- [25 GbE network equipment summary](#)

Introduction

The network is designed to meet the needs of a high performance and scalable cluster, while providing redundancy and access to management capabilities. The architecture is a leaf and spine model that is based on 25 GbE networking technologies. It uses Dell EMC PowerSwitch S5248F-ON switches for the leaves and Dell EMC PowerSwitch Z9432F-ON switches for the spine.

IPv4 is used for the network layer. The design does not support the use of IPv6 for network connectivity.




Figure 3. Cloudera CDP Private Cloud Base cluster logical network layout

Cluster networks

Three networks are used in the cluster:

- Cluster Data network
- Baseboard Management Controller (BMC) or integrated Dell Remote Access Controller (iDRAC) network
- Edge network

The table below describes the networks, their purposes, and their switch connections.

Table 5. CDP Private Cloud Base network definitions

Network	Description	Available services	Connection	Switch
Cluster Data network	This network carries the bulk of the traffic within the cluster. This network is aggregated within each pod, and pods are aggregated into the cluster switch.	The Cloudera Enterprise services are available on this network.	25 GbE	Top of rack (pod) switches and aggregation switches
BMC or iDRAC network	This network connects the BMC or iDRAC ports and the out-of-band management ports of the switches. It is used for hardware provisioning and management. This network is aggregated into a management switch in each rack.	This network provides access to the BMC and iDRAC functionality on the servers. It also provides access to the management ports of the cluster switches.	1 GbE	Dedicated switch per rack
Edge network	This network provides connectivity from the Edge nodes to an existing network, either directly or by the pod or cluster aggregation switches.	SSH access to Edge nodes is available on this network. Other application services may be configured and available.	25 GbE	Direct to the Edge network, or by pod or aggregation switch

NOTE: The CDP Private Cloud Base services do not support multihoming. The HBase, HDFS, HIVE, Impala, and Ozone services are only accessible on the Cluster Data network.

Connectivity between the cluster and existing network infrastructure can be adapted to specific installations. Common scenarios include:

- The Cluster Data network is isolated from any existing network, and access to the cluster is by the Edge network only.
- The Cluster Data network is exposed to an existing network. In this scenario, the Edge network is either unused or is used for application access or ingest processing.

Each network uses a separate VLAN and dedicated components when possible.

Server node connections

Server connections to the network switches for the Data network use Ethernet technology.

Dell Technologies recommends 25 GbE for new deployments on Dell EMC PowerEdge R650 and Dell EMC PowerEdge R750 servers.

Edge nodes and Utility nodes have two additional network connections. These connections facilitate high-performance cluster access between applications running on those nodes, and the optional Edge network. Bonds are typically configured across the two network cards to ensure network resiliency against a NIC or port failure.

Server connections to the BMC network use a single connection from the iDRAC port to an S3100-ON family management switch in each rack.




Figure 4. Dell EMC PowerEdge R650 Infrastructure node network ports




Figure 5. Dell EMC PowerEdge R750 Worker node network ports

Network high availability

The production network should use bonded connections from the hosts to pairs of switches in each pod. Pod-level switches should use redundant connections to switch pairs at the aggregation level. This configuration provides increased bandwidth capacity and allows operation at reduced capacity if a network port, network cable, or switch fails.

Dell Technologies recommends that the pod network uses two S5248F-ON switches that are configured with VLT and Link Aggregation Groups (LAGs). That configuration provides maximum resilience of the networking layer. Host NIC ports are configured in the active/active mode to ensure that LACP runs on any link that is configured in this state. A port in an Active state automatically initiates negotiations with other ports by initiating LACP packets.

Network fabric

The core network fabric is based on Ethernet technology. For clusters larger than a single pod, an aggregation layer is required. The aggregation layer can be implemented at either Layer 2 (L2) or Layer 3 (L3). The choice depends on the initial size and planned scaling. Layer 2 is preferred for lower cost and medium scalability.

Layer 3 aggregation is recommended for:

- Larger initial deployments over 36 nodes
- Deployments where extreme scale-up is planned to exceed 250 nodes
- Instances where the cluster must be co-located with other applications in a different rack

The scalability depends on the switches that are used and the oversubscription ratio, and is summarized in [Cluster node counts](#).

Pod switches

Each pod uses a Dell EMC PowerSwitch S5248F-ON as the first layer switch.

NOTE: The pod switches are often called Top of Rack (ToR) switches. However, this design splits a physical rack from a logical pod.

The 48-port S5248F-ON provides High-density 10/25 GbE ToR server aggregation with scalable L2 and L3 Ethernet switching and standards-based IPv4 and IPv6 features. It features line-rate performance using nonblocking switch fabrics of 2.0 Tbps (4.0 Tbps full-duplex). It is configured with:

- 48 ports of 25 GbE (SFP28)
- 2 ports of 200 GbE (QSFP28-DD)
- 4 ports of 100 GbE (QSFP28)

25 GbE single pod networking equipment shows the single pod network configuration, with Dell EMC PowerSwitch S5248F-ON switches aggregating the pod traffic.




Figure 6. 25 GbE single pod networking equipment

For a single pod, the ToR switches can act as the aggregation layer for the entire cluster. For multiple-pod clusters, a cluster aggregation layer is required. In this architecture, each pod is managed as a separate entity from a switching perspective. The individual pod ToR switches connect only to the spine switch.

In this design, each pod is managed as a separate entity from a switching perspective, and the individual pod switches connect only to the aggregation switch.

Aggregation switches

The Dell EMC PowerSwitch Z9432F-ON is used for both Layer 2 implementations and can be paired with a second Z9432F-ON switch to create a high availability configuration.

For clusters consisting of more than one pod, this design uses the Z9432F-ON for an aggregation switch. The PowerSwitch Z9432F-ON can be used for both Layer 2 and Layer 3 implementations.

The Dell EMC PowerSwitch Z9432F-ON is a multiple-rate 400 GbE, 1U spine switch optimized for high performance, ultralow latency data center requirements. The PowerSwitch Z9432F-ON can provide a cumulative bandwidth of 25.6 Tbps nonblocking (full duplex). Its switching fabric can deliver line-rate performance under full load. It can be configured with up to:

- 32 ports of 400 GbE (QSFP56-DD)
- 128 ports of 100 GbE (using a breakout cable)
- 144 ports of 10/25/50 GbE (using a breakout cable)

The figure below shows the multiple-pod network configuration with Layer 2 networking.

The uplink from each S5248F-ON pod switch to the aggregation layer uses four 100 GbE interfaces in a bonded configuration. This configuration provides a collective bandwidth of 400 Gb from each pod.




Figure 7. Multiple pod networking equipment

Layer 3 cluster aggregation

Layer 3 aggregation is used for large clusters and supports scalability beyond a single set of aggregation switches.

[Multiple pod view with Layer 3 ECMP](#) illustrates cluster aggregation using the PowerSwitch Z9432F-ON and the PowerSwitch S5248F-ON using a leaf-spine model that is based on ECMP and Layer 3 routing.

More details about Layer 3 leaf-spine deployment can be found in the Dell Technologies Info Hub document, [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).




Figure 8. Multiple pod view with Layer 3 ECMP

iDRAC management network

In addition to the Cluster Data network, a separate network is provided for cluster management - the iDRAC (or BMC) network. The iDRAC management ports are all aggregated into a per-rack Dell EMC PowerSwitch S3100-ON family switch. This aggregation provides a dedicated iDRAC or BMC network, for hardware provisioning and management. Switch management ports are also connected to this network.

If out-of-band management is required the management switches can be connected to the core, or connected to a dedicated Management network.

Core network integration

The aggregation layer functions as the network core for the cluster. In most instances, the cluster connects to a larger core within the enterprise, as shown in [Multiple pod networking equipment](#).

With the Dell EMC PowerSwitch Z9432F-ON, two 100 GbE ports are reserved for connection to the core. Details of the connection are site-specific and must be determined as part of the deployment planning.

PowerScale network integration

When PowerScale is used for primary HDFS storage, the Cludera cluster nodes are connected using the same switching architecture that is described above. The PowerScale nodes are connected to their own switching layer and then linked to the aggregation layer.

PowerScale nodes use a dedicated backend network in addition to the Cluster Data network. The backend network is used for internode communication between the PowerScale nodes. This network can use Ethernet or InfiniBand technology.

Every PowerScale node has two 25 GbE connections to the Cluster Data network and two connections to the backend network.

Depending on customer requirements, PowerScale can be networked directly to the leaf topology of the Cludera CDP Private Cloud Base, or connected to the aggregation layer at the spine. See [Cluster sizing and scaling](#) for sizing and scaling information.




Figure 9. PowerScale networking

25 GbE network equipment summary

The number of cables that are needed for a cluster are summarized in:

- Per-rack network equipment
- Per-pod network equipment
- Per-cluster aggregation network switches for multiple pods

Table 6. Per-rack network equipment

Component	Quantity
Total racks	One (12 nodes nominal)
Management switch	One Dell EMC PowerSwitch S3100-ON family
Switch interconnect cable	One 1 GbE cable (to the next rack management switch)

Table 7. Per-pod network equipment

Component	Quantity
Total racks	Three (36 nodes)
ToR pod switches	Two Dell EMC PowerSwitch S5248F-ON
Pod uplink cables	Four 100 Gb QSFP+ cables (to the aggregate switch)

Table 8. Per-cluster aggregation network switches for multiple pods

Component	Quantity
Total Pods	Eight
Aggregation layer switches	Two Dell EMC PowerSwitch Z9432F-ON

See the [Dell EMC PowerSwitch Z9432F-ON Spec Sheet](#) for more information about the Z9432F-ON.

The required cluster networking equipment is summarized in [Per-node network cables required](#).

Table 9. Per-node network cables required

Description	1 GbE cables	25 GbE connections
Master nodes	One per node	Two per node
Edge nodes	One per node	Two per node
Worker nodes	One per node	Two per node

NOTE: 25 GbE node connections typically use a QSFP28 (100 Gbps) to SFP28 (25 Gbps) breakout cable. The cable count is typically one-fourth the number of connections in [Per-node network cables required](#).

Infrastructure configurations

Dell Technologies recommends the following server and storage configurations for use with Cloudera CDP Private Cloud Base. This chapter includes alternative configurations for some nodes and deployments that do not use all these configurations. The recommendations include general guidance on customizing the configurations. See [Cluster scaling](#) for additional information about overall cluster sizing.

Topics:

- [Infrastructure nodes](#)
- [Worker nodes](#)
- [PowerScale node](#)

Infrastructure nodes

Infrastructure nodes are used to host the critical cluster infrastructure services, including:

- NameNode processes
- YARN ResourceManager
- ZooKeeper
- HBase masters
- Cloudera Manager
- Supporting databases

Dell Technologies recommends the configuration that is listed in [Infrastructure nodes configuration](#) as a starting point. This configuration is optimized for reliability, provides high performance, and is consistent with recommendations from Cloudera.

Table 10. Infrastructure nodes configuration

Machine function	Component
Platform	Dell EMC PowerEdge R650 server
Chassis	1U 2.5" chassis with up to 10 hard drives (SAS or SATA), including a maximum of four universal drives, three PCIe slots, and two CPUs
Chassis configuration	Riser configuration 0, two CPUs, half length, low profile, three 16 slots, software GPU capable
Power supply	Mixed mode, 1100 W redundant power supply
Processor	Dual Intel Xeon Gold 6326 2.9 G, 16 C/32 T, 11.2GT/s, 24M cache, turbo, HT (185 W) DDR4-3200
Memory	128 GB - Sixteen 16 GB RDIMM, 3200MT/s, single rank
Persistent memory	None
OCP network card	Intel E810-XXV dual port 10/25 GbE SFP28 OCP NIC 3.0
Extra network card	Intel E810-XXV dual port 10/25 GbE SFP28 PCIe low profile adapter
Storage controller	PERC H745 Front SAS RAID controller
Disk - HDD	None
Disk - SSD	Dual 480 GB SSD SATA mixed use 6 Gbps 512e 2.5" hot-plug S4610 drives

Table 10. Infrastructure nodes configuration (continued)

Machine function	Component
	Six 1.6 TB SSD SAS mixed use 12 Gbps 512e 2.5" hot-plug AG drives, three Drive Writes per Day (DWPD)
Disk - NVMe	One 1.6 TB enterprise NVMe mixed use U2, G4, P5600 with carrier
Boot configuration	From PERC controller

Dell Technologies recommends the disk volume and partition layouts for this set of machines, which are listed in:

- [Infrastructure node volumes](#)
- [Infrastructure node partitions](#)

Table 11. Infrastructure node volumes

Usage	Volume type	Physical disks	Volume Id
Operating system	RAID 1	Two 480 GB SATA SSD	0
HDFS metadata and operational databases	RAID 6	Six 1.6 TB SAS SSD	1
ZooKeeper and NameNode journal	No RAID	One 1.6 TB NVMe	2

Table 12. Infrastructure node partitions

Mount point	Size	File system type	Volume Id	Partition type	Description
/boot	1024 MB	ext4	0	Primary	This partition contains BIOS start-up files that must be within first 2 GB of disk.
/	100 GB	ext4	0	LVM	This partition contains the root file system.
swap	4 GB	swap	0	swap	This partition contains the operating system swap space partition.
/home	1 GB	ext4	0	LVM	This partition contains the user home directories.
/var	~ 350 GB	ext4	0	LVM	This partition contains variable data like system logging files, databases, mail and printer spool directories, transient, and temporary files.
/journal/zookeeper	800 MB	ext4	2	LVM	This partition is used for the ZooKeeper data log. The ZooKeeper configuration property <code>dataLogDir</code> must be changed to match this path at installation time.
/journal/dfs	800 MB	ext4	2	LVM	This partition is used to store the NameNode transactions (edits) files. The configuration property <code>dfs.namenode.edits.dir</code> must be changed to match this path at installation time.
/var/lib/dfs	2 TB	ext4	1	LVM	This partition is used for the NameNode (<code>fsimage</code>) table. The configuration property <code>dfs.namenode.name.dir</code> must be changed to match this path at installation time.
/var/lib/zookeeper	500 MB	ext4	1	LVM	This partition is used for the ZooKeeper database. The ZooKeeper configuration property <code>dataDir</code> must be changed to match this path at installation time.

Table 12. Infrastructure node partitions (continued)

Mount point	Size	File system type	Volume Id	Partition type	Description
<code>/var/lib/pgsql</code>	~ 4 TB	ext4	1	LVM	This partition contains the operational data directory for databases. This directory primarily contains the Cloudera Manager databases, since the PostgreSQL data directory (PGDATA) is typically <code>/var/lib/pgsql</code> . Alternatives to PostgreSQL should be configured to store their datafiles here.

Dell Technologies recommends this Infrastructure node configuration for Master nodes in the cluster. The configuration is sized to support Master nodes in a production deployment.

Edge nodes and Utility nodes should use this configuration as a starting point. You can change the processor, memory, and storage recommendations to specialize those nodes.

The configuration includes four network ports to provide two ports for the Cluster Data network, and two ports for the Edge network or other external connections.

Two SSDs in a RAID 1 configuration are used for the operating system volume. The `swap` partition is small since swapping causes excessive latency for critical cluster infrastructure. The home directories are allocated in a separate small partition since user files should not be stored on infrastructure nodes. Most of the storage is allocated to the `/var` partition for runtime files. You can use LVM to adjust the storage allocation between `/`, `/home`, and `/var` for specific needs.

A six SSD RAID 6 volume is used for most of the Infrastructure node storage. This configuration provides a good balance between performance, data durability, storage efficiency, and administration overhead. This volume is divided into partitions for:

- NameNode file system data
- ZooKeeper data
- Cloudera Manager database storage
- Hive metastore
- Ranger database
- Any other required operational databases

You can use LVM to adjust the storage allocation for specific needs.

An alternative configuration is to set up a four SSD RAID 10 volume for database storage, and a two SSD RAID 1 for NameNode and ZooKeeper data. This configuration option provides slightly better performance for database writes, but adds administration overhead and increases recovery time if drives fail.

A single NVMe drive is used for storage of ZooKeeper and quorum journals.

Worker nodes

Worker nodes provide compute and storage resources for the workloads running on the cluster. They can be optimized for storage-dense, compute-dense, or mixed workloads depending on the intended workloads and services running on them.

CDP Private Cloud Base supports various hybrid solutions where compute tasks are separated from data storage and where data can be accessed from remote clusters. This architecture provides the following alternative Worker node configurations:

- [General purpose Worker nodes](#)
- [All flash Worker nodes](#)
- [GPU accelerated Worker nodes](#)
- [PowerScale Worker nodes](#)

General purpose Worker nodes

Dell Technologies recommends the General purpose Worker node configuration as a starting point for most clusters. This configuration is optimized for reliability, provides high performance, and is consistent with recommendations from Cloudera.

Table 13. General purpose Worker node configuration

Machine function	Component
Platform	Dell EMC PowerEdge R750 server
Chassis	2U 3.5" chassis with up to 12 HDDs (SAS or SATA), two 2.5" rear HDDs (SAS or SATA), adapter PERC
Chassis configuration	Riser configuration 6, four 8, two 16 slots
Power supply	Dual hot-plug, 1400 W, mixed mode, redundant (1+1) power supplies
Processor	Dual Intel Xeon Gold 6348 2.6 G, 28 C/56 T, 11.2 GT/s, 42 M cache, turbo, HT (235 W) DDR4-3200
Memory	512 GB - Sixteen 32 GB RDIMM, 3200 MT/s, dual rank
Persistent memory	None
OCP network card	Intel E810-XXV dual port 10/25 GbE SFP28 OCP NIC 3.0
Extra network card	None
Storage controller	Dell HBA355i Adapter RAID controller, full height
Disk - HDD	Twelve 4 TB 7.2 K RPM NLSAS 12 Gbps 512n 3.5" hot-plug hard drives
Disk - SSD	None
Disk - NVMe	3.2 TB enterprise NVMe mixed use, U2, G4, P5600 FlexBay
Boot configuration	BOSS-S2 controller card with two M.2 480 GB (RAID 1)

Dell Technologies recommends the disk volume and partition layouts for this set of machines, which are listed in:

- [General purpose Worker node volume layout](#)
- [General purpose Worker node file system layout](#)

Table 14. General purpose Worker node volume layout

Usage	Volume type	Physical disks	Volume Id
Operating system	RAID 1	Two 480 GB M.2 SSD	0
HDFS data	No RAID	Twelve 4 TB NL-SAS	1-12
Temporary data	No RAID	One 3.2 TB NVMe	13

Table 15. General purpose Worker node file system layout

Mount point	Size	File system type	Volume Id	Partition type	Description
boot	1024 MB	ext4	0	Primary	This partition contains BIOS start-up files that must be within first 2 GB of disk.
/	100 GB	ext4	0	LVM	This partition contains the root file system.
swap	4 GB	swap	0	swap	This partition contains the operating system swap space partition.

Table 15. General purpose Worker node file system layout (continued)

Mount point	Size	File system type	Volume Id	Partition type	Description
/home	1 GB	ext4	0	LVM	This partition contains the user home directories.
/var	~350 GB	ext4	0	LVM	This partition contains variable data like system logging files, databases, mail and printer spool directories, transient, and temporary files.
/data/<n>	4 TB	ext4	1-12	Primary	This partition is used for HDFS data as 12 individual file systems.
/data/ssd	3.2 TB	ext4	13	Primary	This partition is used for temporary files.

The General purpose Worker node configuration is sized for a typical mix of storage and computes in a Cloudera CDP Private Cloud Base cluster.

Two network ports are included for connection to the Cluster Data network.

Two SSDs in a RAID 1 configuration using the boot optimized storage system (BOSS) are used for the operating system volume. The `swap` partition is small since swapping causes excessive latency for running jobs. The home directories are allocated in a separate, small partition since user files should not be stored on Worker nodes. Most of the storage is allocated to the `/var` partition for runtime files. You can use LVM to adjust the storage allocation between `/`, `/home`, and `/var` for specific needs.

Twelve 3.5" NL-SAS hard drives are used for the primary data storage. These drives are mounted as individual partitions and provide approximately 48 TB of raw storage. Cloudera supports a maximum of 100 TB per node for HDFS storage. You can increase the drive size to 8 TB for maximum storage capacity, but doing so has a negative impact on cluster performance. A larger HDFS storage capacity increases the time that is required for background scans and block reports. It also increases the recovery time should a node fail.

A single NVMe drive is used for storage of temporary files such as MapReduce temporary and spill files, and Spark cache. You can also use this drive for HBase tiered cache or tiered HDFS storage. You can also increase the size of this drive or add an additional U.2 NVMe drive as necessary.

The recommended memory size of 512 GB is intended for Worker nodes that run jobs which can benefit from additional memory such as Spark and HBase region servers. You can reduce the memory allocation for nodes that primarily provide storage services with minimal compute capability.

All flash Worker nodes

Dell Technologies recommends the All flash configuration for Worker nodes that require high-performance storage. This configuration is optimized for reliability, provides high performance, and is consistent with recommendations from Cloudera.

Table 16. All flash Worker node configuration

Machine function	Component
Platform	Dell EMC PowerEdge R750 server
Chassis	2U 3.5" chassis with up to 12 HDDs (SAS or SATA), two 2.5" rear HDDs (SAS or SATA), adapter PERC
Chassis configuration	Riser configuration 6, four 8, two 16 slots

Table 16. All flash Worker node configuration (continued)

Machine function	Component
Power supply	Dual hot-plug, 1400 W, mixed mode, redundant (1+1) power supplies
Processor	Dual Intel Xeon Gold 6348 2.6 G, 28 C/56 T, 11.2 GT/s, 42 M cache, turbo, HT (235 W) DDR4-3200
Memory	512 GB - Sixteen 32 GB RDIMM, 3200 MT/s, dual rank
Persistent memory	None
OCP network card	Intel E810-XXV dual port 10/25 GbE SFP28 OCP NIC 3.0
Extra network card	None
Storage controller	Dell HBA355i Adapter RAID controller, full height
Disk - HDD	None
Disk - SSD	Twelve 3.84 TB SSD vSAS mixed use 12 Gbps 512e 2.5" with 3.5" HYB CARR hot-plug AG drives, three DWPD
Disk - NVMe	3.2 TB enterprise NVMe mixed use, U2, G4, P5600 FlexBay
Boot configuration	BOSS-S2 controller card with two M.2 480 GB (RAID 1)

Dell Technologies recommends the disk volume and partition layouts for this set of machines, which are listed in:

- [All flash Worker node volume layout](#)
- [All flash Worker node file system layout](#)

Table 17. All flash Worker node volume layout

Usage	Volume type	Physical disks	Volume Id
Operating system	RAID 1	Two 480 GB M.2 SSD	0
HDFS data	No RAID	Twelve 3.84 TB vSAS	1-12
Temporary data	No RAID	One 3.2 TB NVMe	13

Table 18. All flash Worker node file system layout

Mount point	Size	File system type	Volume Id	Partition type	Description
boot	1024 MB	ext4	0	Primary	This partition contains BIOS start-up files that must be within first 2 GB of disk.
/	100 GB	ext4	0	LVM	This partition contains the root file system.
swap	4 GB	swap	0	swap	This partition contains the operating system swap space partition.
/home	1 GB	ext4	0	LVM	This partition contains the user home directories.
/var	~350 GB	ext4	0	LVM	This partition contains variable data like system logging files,

Table 18. All flash Worker node file system layout (continued)

Mount point	Size	File system type	Volume Id	Partition type	Description
					databases, mail and printer spool directories, transient, and temporary files.
/data/<n>	3.84 TB	ext4	1-12	Primary	This partition is used for HDFS data as 12 individual file systems.
/data/ssd	3.2 TB	ext4	13	Primary	This partition is used for temporary files.

The All flash Worker node configuration is sized for a typical mix of storage and computes in a Cloudera CDP Private Cloud Base cluster. It is almost the same as the General purpose Worker node configuration except that it uses flash media for the primary storage. This configuration provides high-performance storage for workloads that can benefit from it. Worker nodes running HBase region servers or processing data pipelines with intermediate results benefit from this configuration.

Twelve 3.5" vSAS SSDs are used for the primary data storage. These drives are mounted as individual partitions and provide approximately 46 TB of raw storage.

GPU accelerated Worker nodes

Dell Technologies recommends the GPU accelerated configuration for Worker nodes that require a GPU accelerator. This configuration is optimized for reliability, provides high performance, and is consistent with recommendations from Cloudera.

Table 19. GPU accelerated Worker node configuration

Machine function	Component
Platform	Dell EMC PowerEdge R750 server
Chassis	2U 2.5" chassis with up to 16 SAS or SATA drives
Chassis configuration	Riser configuration 2, full length, four 16, two 8 slots, DW GPU capable
Power supply	Dual, hot-plug, 2400 W, D mixed mode, redundant power supplies
Processor	Dual Intel Xeon Gold 6330 2 G, 28 C/56 T, 11.2 GT/s, 42 M cache, turbo, HT (205 W) DDR4-2933
Memory	512 GB - Sixteen 32 GB RDIMM, 3200 MT/s, dual rank
Persistent memory	None
OCP network card	Intel E810-XXV dual port 10/25 GbE SFP28 OCP NIC 3.0
Extra network card	None
Storage controller	Dell HBA355i Front RAID controller
Disk - HDD	Sixteen 2.4 TB 10 K RPM SAS 12 Gbps 512e 2.5" hot-plug hard drives
Disk - SSD	None
Disk - NVMe	6.4 TB enterprise NVMe mixed use, AG drive, AIC, PCIe 4.0, DIB
Boot configuration	BOSS-S2 controller card with two M.2 480 GB (RAID 1)

Table 19. GPU accelerated Worker node configuration (continued)

Machine function	Component
GPU	NVIDIA Ampere A30, PCIe, 165 W, 24 GB passive, double wide, full height GPU

Dell Technologies recommends the disk volume and partition layouts for this set of machines, which are listed in:

- [GPU accelerated Worker node volume layout](#)
- [GPU accelerated Worker node file system layout](#)

Table 20. GPU accelerated Worker node volume layout

Usage	Volume Type	Physical disks	Volume Id
Operating system	RAID 1	Two 480 GB M.2 SSD	0
HDFS data	No RAID	Sixteen 2.4 TB SAS	1-16
Temporary data	No RAID	One 6.4 TB NVMe	17

Table 21. GPU accelerated Worker node file system layout

Mount point	Size	File system type	Volume Id	Partition type	Description
boot	1024 MB	ext4	0	Primary	This partition contains BIOS start-up files that must be within first 2 GB of disk.
/	100 GB	ext4	0	LVM	This partition contains the root file system.
swap	4 GB	swap	0	swap	This partition contains the operating system swap space partition.
/home	1 GB	ext4	0	LVM	This partition contains the user home directories.
/var	~350 GB	ext4	0	LVM	This partition contains variable data like system logging files, databases, mail and printer spool directories, transient, and temporary files.
/data/<n>	2.4 TB	ext4	1-16	Primary	This partition is used for HDFS data as 16 individual file systems.
/data/ssd	6.4 TB	ext4	17	Primary	This partition is used for temporary files.

The GPU accelerated Worker node configuration is sized for a typical mix of storage and computes in a Cloudera CDP Private Cloud Base cluster. It is based on the General purpose Worker node configuration and adds GPU acceleration. Worker nodes running Spark jobs using the NVIDIA RAPIDS Accelerator for Apache Spark can benefit from the accelerators.

Sixteen 2.5" SAS hard drives are used for the primary data storage. These drives are mounted as individual partitions and provide approximately 38 TB of raw storage.

A single NVMe drive is used for storage of temporary files such as MapReduce temporary and spill files, and Spark cache. This drive is sized to provide extra storage space compared to the General purpose Worker node. You can allocate the extra space to HDFS and use file or directory storage policies to allocate data to the NVME storage. This strategy can be useful for GPU processing since disk I/O can limit GPU performance.

The baseline recommendation uses NVIDIA A30 GPUs since the expected usage in a Worker node is acceleration of Spark ETL-style processing. The NVIDIA A100 can be used as an alternative for compute-intensive jobs such as machine learning training. You can configure these nodes with one or two GPUs, depending on workload requirements.

PowerScale Worker nodes

Dell Technologies recommends the PowerScale Worker node configuration for Worker nodes in configurations where PowerScale is used as the primary HDFS storage. This configuration is optimized for reliability, provides high performance, and is consistent with recommendations from Cloudera.

Table 22. PowerScale Worker node configuration

Machine function	Component
Platform	Dell EMC PowerEdge R650 server
Chassis	1U 2.5" chassis with up to 10 NVMe drives, three PCIe slots, two CPUs
Chassis configuration	Riser configuration 4, half length, low profile, three 16 slots with R1 paddle, SW GPU capable
Power supply	Dual hot-plug, 1400 W, mixed mode, redundant (1+1) power supplies
Processor	Dual Intel Xeon Gold 6330 2 G, 28 C/56 T, 11.2 GT/s, 42 M cache, turbo, HT (205 W) DDR4-2933
Memory	512 GB - Sixteen 32 GB RDIMM, 3200 MT/s, dual rank
Persistent memory	None
OCP network card	Intel E810-XXV dual port 10/25 GbE SFP28 OCP NIC 3.0
Extra network card	None
Storage controller	None
Disk - HDD	None
Disk - SSD	None
Disk - NVMe	3.2 TB enterprise NVMe mixed use, U2, G4, P5600 with carrier
Boot configuration	BOSS-S2 controller card with two M.2 480 GB (RAID 1)

Dell Technologies recommends the disk volume and partition layouts for this set of machines, which are listed in:

- [PowerScale Worker node volume layout](#)
- [PowerScale Worker node file system layout](#)

Table 23. PowerScale Worker node volume layout

Usage	Volume type	Physical disks	Volume Id
Operating system	RAID 1	Two 480 GB M.2 SSD	0
Temporary data	No RAID	One 3.2 TB NVMe	1

Table 24. PowerScale Worker node file system layout

Mount point	Size	File system type	Volume Id	Partition type	Description
boot	1024 MB	ext4	0	Primary	This partition contains BIOS start-up files that must be within first 2 GB of disk.
/	100 GB	ext4	0	LVM	This partition contains the root file system.
swap	4 GB	swap	0	swap	This partition contains the operating system swap space partition.
/home	1 GB	ext4	0	LVM	This partition contains the user home directories.
/var	~350 GB	ext4	0	LVM	This partition contains variable data like system logging files, databases, mail and printer spool directories, transient, and temporary files.
/data/ssd<n>	3.2 TB	ext4	1	Primary	This partition is used for temporary files.

The PowerScale Worker node is intended for use in configurations where PowerScale provides the primary HDFS storage. This configuration does not include any drives for HDFS storage.

Two network ports are included for connection to the Cluster Data network.

Two NVMe drives are used for storage of temporary files such as MapReduce temporary and spill files, and Spark cache. The amount of required local storage is based on the workloads running and the amount of local storage they need. For general guidance, the temporary storage that is required should be between 10% and 20% of the overall cluster HDFS storage. It should be distributed across all the Worker nodes. Between 5 TB and 10 TB per Worker node is appropriate for most clusters using Dell EMC PowerScale Isilon H5600 storage nodes. You can change both the size, and number of NVMe drives to meet these requirements. For more details about sizing PowerScale-based clusters, see [Cluster scaling](#).

PowerScale node

Dell Technologies recommends the PowerScale Isilon configuration for storage in clusters using PowerScale for their primary HDFS storage.

Table 25. PowerScale Isilon configuration

Machine function	Component
Model	Dell EMC PowerScale Isilon H5600 (hybrid)
Chassis	4U
Nodes per chassis	4U
Node storage	Twenty 10 TB 3.5" 4 kn SATA HDD

Table 25. PowerScale Isilon configuration (continued)

Machine function	Component
Node cache	Two 3.2 TB SSD
Usable capacity per chassis	500 TB
Front-end networking	Two 25 GbE (SFP28)
Infrastructure (back-end) networking	Two InfiniBand QDR or two 40 GbE (QSFP+)
Operating system	OneFS 9.2

The PowerScale Isilon configuration is sized for typical usage as primary HDFS storage with Cloudera CDP Private Cloud Base clusters and the PowerScale Worker node configuration.

Two Ethernet network ports per node are included for connection to the Cluster Data network or a PowerScale storage network. Two additional network ports are included for connection to the PowerScale back-end network. These additional ports can be either InfiniBand QDR or 40 GbE, depending upon on-site preferences.

One Dell EMC PowerScale Isilon H5600 chassis supports four Dell EMC PowerScale Isilon H5600 nodes. This configuration provides approximately 594 TB of usable storage. At 85% utilization, 500 TB of HDFS storage is a good guideline for available storage per chassis.

This configuration assumes that the PowerScale nodes are primarily used for HDFS storage. If the PowerScale nodes are used for other storage applications or clusters, you must account for it in the overall cluster sizing. You can also use other Dell EMC PowerScale Isilon H5600 drive configurations. For more details about sizing PowerScale-based clusters, see [Cluster scaling](#).

Software infrastructure

The Dell Technologies Validated Design for Cloudera CDP Private Cloud Base requires specific versions of Dell and third-party software.

Topics:

- [Software components](#)

Software components

The software components and versions that are validated for Cloudera CDP Private Cloud Base are listed in the table below.

Table 26. Software components

Category	Component	Version
Operating system	Red Hat Enterprise Linux Server	8.2
	Linux kernel	4.18.0-193.19.1.el8_2.x86_64
File system	XFS	N/A
Java Virtual Machine	Open JDK	8 (8u232-b09)
Cloudera Data Platform	CDP Private Cloud Base	7.1.7
	Cloudera Manager	7.4.4
Firmware	iDRAC	4.40.20.00
	NVIDIA Ampere A30	470.57.02
	Intel E810-XXV	20.0.18
	PERC H745	51.14.0-3900
	Dell HBA355i	14.25.12.00
GPU computing platform	NVIDIA CUDA	11.4.2
Database	PostgreSQL	10.15-1

Validation

Once the CDP Private Cloud Base environment has been deployed, it must be validated to ensure that it is configured correctly and performs successfully.

Topics:

- [Overview](#)
- [Component validation](#)
- [Findings](#)
- [Summary](#)

Overview

This chapter describes functional validation of multiple Cloudera services that Dell Technologies deployed as a part of the Cloudera CDP Private Cloud Base software stack. This testing begins with services health check up to ensure that all the Cloudera services are running in good health.

Dell Technologies has also tested integrating a Dell EMC PowerScale storage array with Cloudera CDP Private Cloud Base 7.1.7. This platform provides Hadoop clients with direct access to big data through the HDFS interface. The Hadoop NameNode functionality that is integrated into PowerScale simplifies deploying and scaling of CDP Private Cloud Base. Dedicated NameNode servers are not required.

Testing also covers running sample workloads on various components of Cloudera CDP Private Cloud Base, including:

- HDFS
- Hive
- HBase
- Spark
- MapReduce
- Spark and Hive integration
- Hue
- Spark GPU test

Component validation

The testing objective was to run sample workloads to validate the functionality of different components and services that Dell Technologies installed as parts of Cloudera CDP Private Cloud Base.

TeraSuite test

About this task

The TeraSuite workload tool combines testing of HDFS and MapReduce layers of a Hadoop cluster. Its goal is to generate, sort, and validate a configurable amount of data as fast as possible. This test is designed to exercise the compute and local storage configurations with concurrent HDFS access.

Dell Technologies ran the TeraSuite commands for 1 TB of data in each test case. The sum of TeraGen, TeraSort, and TeraValidate execution times determined the ranking. Lower scores are better.

Steps

1. This example `teragen` command generates data:

```
time hadoop jar hadoop-mapreduce-examples-3.1.1.7.1.7.0-551.jar teragen \  
-Ddfs.blocksize=536870912 -Dmapreduce.job.maps=240 -Dmapreduce.job.reduces=120 \  
-Dmapreduce.map.speculative=true -Dmapreduce.map.output.compress=true 10000000000 \  
hdfs://pvcmaster1.orange.local:8020/user/root/teragen1
```

2. This example `terasort` command sorts the generated data:

```
time hadoop jar hadoop-mapreduce-examples-3.1.1.7.1.7.0-551.jar terasort \  
-Ddfs.blocksize=536870912 -Dmapreduce.job.maps=240 -Dmapreduce.job.reduces=120 \  
-Dmapreduce.map.speculative=true -Dmapreduce.map.output.compress=true \  
/user/root/teragen1 /user/root/terasort1
```

3. This example `teravalidate` command validates that the generated data has been sorted:

```
time hadoop jar hadoop-mapreduce-examples-3.1.1.7.1.7.0-551.jar teravalidate \  
-Ddfs.blocksize=536870912 -Dmapreduce.job.maps=240 -Dmapreduce.job.reduces=120 \  
-Dmapreduce.map.speculative=true -Dmapreduce.map.output.compress=true \  
/user/root/terasort1 /user/root/teravalidate
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the cluster HDFS and MapReduce layers.

DSFIO test

About this task

The TestDFSIO workload tool is a read and write test for HDFS that is included with CDP. TestDFSIO was run with many files to create multiple execution threads.

Steps

1. This example `TestDFSIO` command performs a TestDFSIO write test:

```
yarn jar /opt/cloudera/parcels/CDH/jars/hadoop-mapreduce-client \  
-jobclient-3.1.1.7.1.7.0-551-tests.jar TestDFSIO -write -nrFiles 5000 -size 128MB
```

2. This example `TestDFSIO` command performs a TestDFSIO read test:

```
yarn jar /opt/cloudera/parcels/CDH/jars/hadoop-mapreduce-client \  
-jobclient-3.1.1.7.1.7.0-551-tests.jar TestDFSIO -read -nrFiles 5000 -size 128MB
```

Results

At the conclusion of this test, Dell Technologies proved the validity of HDFS reads and writes.

MapReduce test

About this task

Hadoop MapReduce is a programming model that is used to process bulk data. Programs for MapReduce can be run in parallel. They deliver high-performance, large-scale data analyses in the cluster.

Step

- This example python word count program tests the functional validation of the MapReduce service:

```
hadoop jar /opt/cloudera/parcels/CDH/jars/hadoop-streaming-3.1.1.7.1.7.0-551.jar \  
-file /root/final_test/mapper.py -mapper /root/final_test/mapper.py -file \  
/root/final_test/reducer.py -reducer /root/final_test/reducer.py -input \  
/user/root/words.txt -output /user/root/mp_reduce"
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the MapReduce service.

Spark test

About this test

Apache Spark achieves high performance for both batch and streaming data using:

- A state-of-the-art Directed Acyclic Graph (DAG) scheduler
- A query optimizer
- A physical execution engine

Step

- This example sparkpi application tests the functional validation of the Spark service:

```
spark-submit --class org.apache.spark.examples.SparkPi --master yarn \  
/opt/cloudera/parcels/CDH/jars/spark-examples_2.11-2.4.7.7.1.7.0-551.jar 10
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the Spark service .

Spark GPU test

About this task

Cloudera Data Science (CDS) 3.1 for GPUs is an add-on service. It enables you to take advantage of the RAPIDS Accelerator for Apache Spark to accelerate Spark 3 performance on existing CDP Private Cloud Base clusters.

Step

- This example Spark shell command tests the functional validation of the Spark GPU service (CDS 3.1).

```
spark3-shell --master yarn --conf spark.task.resource.gpu.amount=1 --conf  
spark.rapids.sql.concurrentGpuTasks=1 \  
--conf spark.sql.files.maxPartitionBytes=256m --conf spark.locality.wait=0s --conf  
spark.sql.adaptive.enabled=true \  
--conf spark.rapids.memory.pinnedPool.size=2G --conf "spark.rapids.sql.enabled=true"  
--conf \  
"spark.executor.memoryOverhead=5g" --conf  
spark.sql.adaptive.advisoryPartitionSizeInBytes=1
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the Spark GPU service (CDS 3.1).

Hive test

About this task

Hive is a data warehouse software project that is built on top of Hadoop to provide data query and analysis. Hive provides a SQL-like interface to query data that is stored in various databases and file systems that integrate with Hadoop.

These example queries test the functional validation of the Hive service.

Steps

1. Connect to Hive from the Beeline thin client:

```
!connect jdbc:hive//<Namenode>:10000/default
```

2. Create a database named TEST:

```
CREATE DATABASE TEST
```

3. Create a table:

```
CREATE TABLE TEST.Sales_Data(StoreLocation VARCHAR(30),Product VARCHAR(30),\
OrderDate DATE,Revenue DECIMAL(10,2))
```

4. Insert data into the table:

```
Insert into Sales_Data Values('Bangalore','Nutella','2021-07-16',7455.67),\
('Bangalore','Peanut Butter','2021-07-16',5316.89),('Bangalore','Milk','2021-07-16',\
2433.76),('Hyderabad','Bananas','2021-07-16',9456.01),('Hyderabad','Nutella',\
'2021-07-16',3644.33),('Hyderabad','Peanut Butter', '2021-07-16', 8988.64),\
('Hyderabad','Milk','2021-07-16', 1621.58)
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the Hive service.

HBase test

About this task

HBase is a column-oriented, nonrelational database management system that runs on top of HDFS. HBase provides a fault-tolerant way to store sparse datasets, which are common in many big data use cases.

These example HBase queries were used to test the functional validation of the HBase service.

Steps

1. Create a table named **history** with two columns; **home** and **away**:

```
create 'history', 'home', 'away'
```

2. Insert one row into the table:

```
put 'history','1','home data:name','jim'
put 'history','row1','home:city','Boston'
```

3. Delete the table:

```
drop 'history'
```

Results

At the conclusion of this test, Dell Technologies proved the validity of the HBase service.

Findings

All the services running inside the Cloudera CDP Private Cloud Base cluster must be healthy, as they have dependencies orders.

For example, MapReduce and YARN have dependencies on HDFS. HDFS must be in a healthy state in order to start the MapReduce and YARN Services.

Summary

All the tests ran successfully, with effective utilization of compute and memory resources. The objective of these tests was not to measure any performance metric data. They were meant to confirm that all the services deployed in the Cloudera CDP Private Cloud Base cluster were functioning correctly.

Summary

Cloudera Data Platform (CDP) Private Cloud Base is the first on-premises release of CDP that combines the former Cloudera and Hortonworks software into a single, comprehensive data platform. CDP Private Cloud Base is also used with CDP Private Cloud Data Services to form CDP Private Cloud.

Topics:

- [Overview](#)
- [We value your feedback](#)

Overview

This document provides design guidance for data analytics infrastructure managers and architects by describing a predesigned, validated, and scalable reference architecture for running CDP Private Cloud Base on Dell EMC hardware infrastructure. Topics that are discussed include:

- The software infrastructure components and versions that were used for CDP Private Cloud Base
- The cluster architecture that was designed for this application, including cluster node definitions, roles, and assignments
- The cluster physical and logical network designs
- Cluster sizing and scaling guidance
- High availability considerations
- Details of the PowerEdge server and PowerSwitch networking configurations

Dell Technologies and Cloudera have been collaborating for over eight years to provide customers with guidance on optimal hardware to streamline the design, planning, and configuration of their Cloudera deployments. Dell Technologies is a Platinum member of the Cloudera IHV Program, the highest level of partnership that indicates ongoing commitments to both Cloudera and customers. This document is based on the collective experience of both companies in deploying and running enterprise production environments for Cloudera software on Dell EMC hardware infrastructure.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Authors: Dell Technologies Integrated Solutions Engineering and Technical Marketing teams

 **NOTE:** For links to additional documentation for this solution, see the [Dell Technologies Info Hub for Data Analytics](#).

References

Additional information can be obtained at the [Dell Technologies Info Hub for Data Analytics](#). If you need additional services or implementation help, contact your Dell Technologies sales representative.

Topics:

- [Dell Technologies documentation](#)
- [Cloudera documentation](#)
- [Services engagements](#)
- [Dell Technologies Customer Solution Centers](#)
- [Dell EMC OpenManage Enterprise Power Manager](#)
- [Dell Technologies Info Hub](#)
- [More information](#)

Dell Technologies documentation

The following Dell Technologies documentation provides additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

Table 27. Dell Technologies documentation

Document type	Location
Server spec sheets	Dell EMC PowerEdge R650 Spec Sheet
	Dell EMC PowerEdge R750 Spec Sheet
Storage spec sheets	Dell EMC Isilon H5600 Spec Sheet
Switch spec sheets	Dell EMC PowerSwitch S3100-ON Series Spec Sheet
	Dell EMC PowerSwitch S5200-ON Series Spec Sheet
	Dell EMC PowerSwitch Z9432F-ON Spec Sheet
	Dell EMC PowerSwitch Z9264F-ON Spec Sheet
Server manuals	Dell EMC PowerEdge R650 Manuals and Documents
	Dell EMC PowerEdge R750 Manuals and Documents
Storage manuals	Dell EMC Isilon H5600 Manuals and Documents
Switch manuals	Dell EMC PowerSwitch S3100-ON Manuals & Documents
	Dell EMC PowerSwitch S5200-ON Series Manuals & Documentation
	Dell EMC PowerSwitch Z9432F-ON Manuals & Documentation
	Dell EMC PowerSwitch Z9264F-ON Manuals & Documentation

Cloudera documentation

The following documentation on the [Cloudera documentation website](#) provides additional and relevant information:

Table 28. Cloudera documentation

Document type	Location
New features in CDP Private Cloud Base	CDP Private Cloud Base 7.1.7 Release Summary
Basic requirements for CDP Private Cloud Base	CDP Private Cloud Base Requirements and Supported Versions
Overview of CDP Private Cloud	CDP Private Cloud Overview
Overview of basic Cloudera Manager architecture and specialized terms	Cloudera Manager Architecture
Overview of CDP Private Cloud Data Services	CDP Private Cloud Data Services documentation
Guided upgrade journey	Cloudera Upgrade Guide Companion
Supported upgrade paths for CDH, HDP, and CDP Private Cloud Base	Supported Upgrade Paths
Upgrade and migration using Replication Manager	Replication Manager considerations while using CDH clusters
Upgrades from Apache Ambari and HDP	Upgrading Ambari and HDP to CDP Private Cloud Base
Node role assignments for planning	Runtime Cluster Hosts and Role Assignments
Overview of network and security	Networking and Security Requirements

Services engagements

Accelerate your Cloudera CDP Data and Analytics Initiatives with ProConsult Advisory Service.

Create a strategy and roadmap for analytics with your data platform, including real-time stream processing and data management, and high performance analytics [ProConsult Advisory Service](#).

Experience the value of an end-to-end assessment that delivers actionable outcomes with ProConsult Advisory Services. Dell Technologies Consulting can help you:

- Secure participation and support from key stakeholders and aligned cross-functional teams sharing a single vision and guiding principles.
- Leverage our experiences with what "good" looks like, including best practice guidance.
- Prepare to immediately execute on findings and recommendations.

Our Cloudera certified experts help customers gain additional value from their Cloudera CDP investments with packaged and customer services such as: Data Analytics IoT Technical Advisory and Implementation services, as well as Hadoop migration services.

[Learn more](#) about Dell Technologies ProConsult Services.

[Contact](#) a Dell Technologies Consulting Services Expert today.

Dell Technologies Customer Solution Centers

Our global network of dedicated [Dell Technologies Customer Solution Centers](#) are trusted environments where world class IT experts collaborate with customers and prospects to share best practices; facilitate in-depth discussions of effective business strategies using briefings, workshops, or Proofs of Concept (PoCs); and help business become more successful and competitive.

Dell Technologies Customer Solution Centers reduce the risks that are associated with new technology investments, and can help improve speed of implementation.

All of the services of the Customer Solution Centers are available to all Dell Technologies customers at no charge. Contact your account team today to submit an engagement request.

Dell EMC OpenManage Enterprise Power Manager

[OpenManage Enterprise Power Manager](#) enables you to optimize the control of server power consumption and mitigate operational risks.

Enhance your visibility into server power consumption and control utilization through the OpenManage Enterprise Power Manager console. This powerful console is integrated with the iDRAC that is embedded in Dell EMC PowerEdge servers. Set policy-based controls to maximize resource use and throttle-back power when performance demand ebbs. By using predefined power policies, OpenManage Enterprise Power Manager can help mitigate operational risks and ensure that your servers and their key workloads continue to operate.

Dell Technologies Info Hub

The [Dell Technologies Info Hub](#) is your one-stop destination for the latest information about Dell EMC Solutions and Networking products. New material is frequently added, so browse often to keep up to date on the expanding Dell Technologies portfolio of cutting-edge products and solutions.

More information

For more information, email ai.assist@dell.com, or contact your Dell Technologies or authorized partner sales representative.