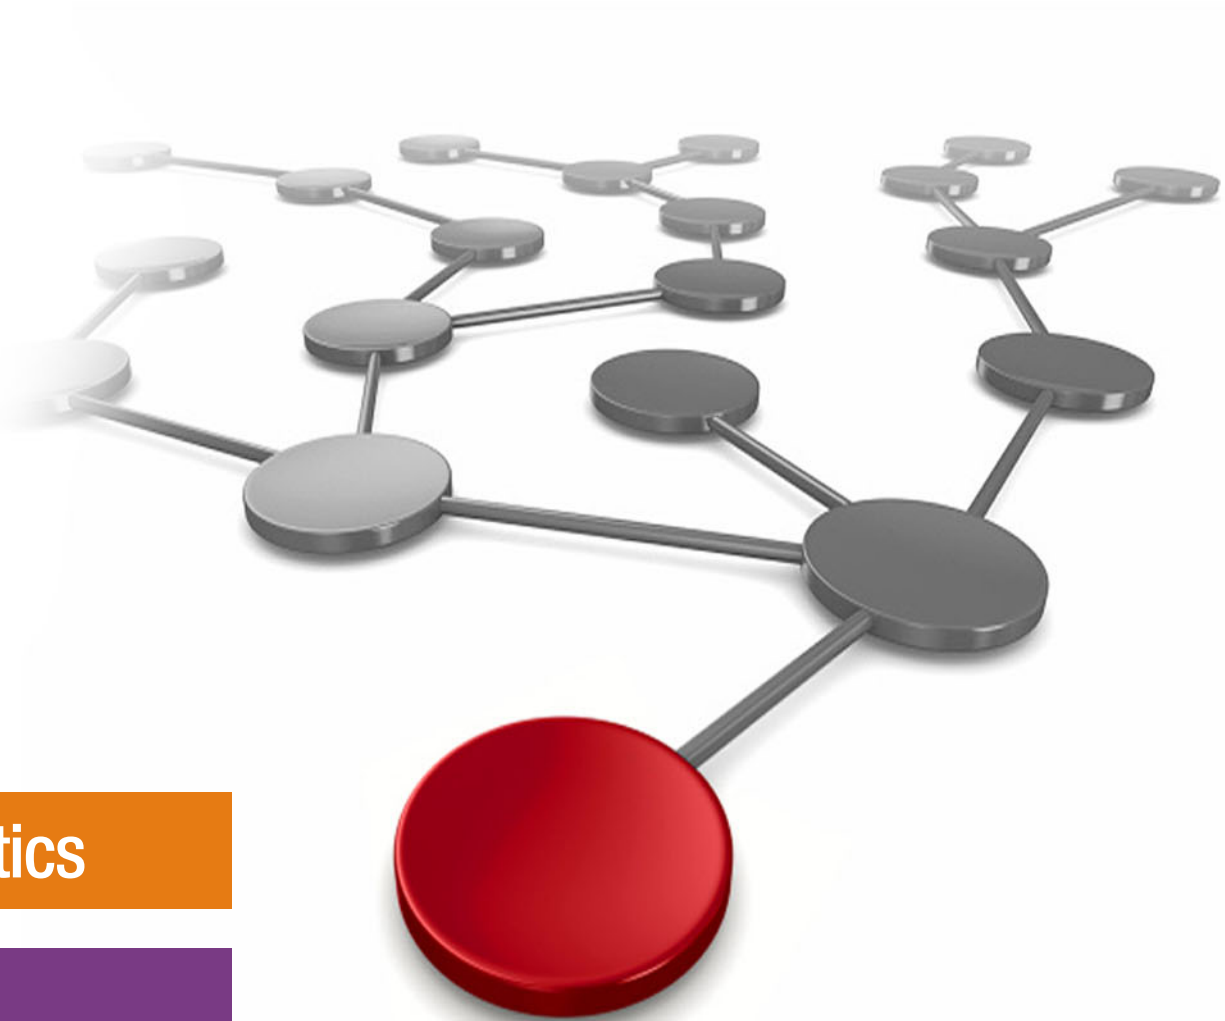# Cloudera Data Platform Private Cloud Base with IBM Spectrum Scale

Wei Gong

Linda Cham

Prashanth Shetty

John Sing

Analytics

Storage

Summary of changes

This section describes the technical changes made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for Cloudera Data Platform Private Cloud Base with IBM Spectrum Scale
as created or updated on August 27, 2021.

August 2021, Minor updates

This revision includes the following new and changed information.

New and Changed information highlights

► Updated CES protocol support in Hadoop environment. See "CES HDFS enabled with other protocol services recommendations" on page 6.
► Updated links from IBM Knowledge Center to IBM Documentation.

April 2021, Minor updates

This revision includes the following new and changed information.

New and Changed information highlights

► Updated support for Data encryption at rest and in transit:
  – "Data Encryption at rest" on page 10.
  – "Transport Layer Security/Secure Sockets Layer encryption" on page 11.

March 2021, Minor updates

This revision includes the following new and changed information.

New and Changed information highlights

► Added support for non-HA NameNode and collocation of Hadoop services on the DataNode. Refer to the following sections:
  – "Alternative cluster configuration" on page 19
  – "DataNode collocation configuration" on page 20
  – "Non-HA NameNode configuration" on page 20
► Updated the following figures:
  Figure 14 on page 18
  Figure 15 on page 19
  Figure 16 on page 20
  Figure 22 on page 26

► Minor updates denoted with change bars

IBM

Cloudera Data Platform Private Cloud Base with IBM Spectrum Scale

This IBM® Redpaper publication provides guidance on building an enterprise-grade data lake by using IBM Spectrum® Scale and Cloudera Data Platform (CDP) Private Cloud Base for performing in-place Cloudera Hadoop or Cloudera Spark-based analytics. It also covers the benefits of the integrated solution and gives guidance about the types of deployment models and considerations during the implementation of these models.

> **Note:** In January 2019, the Cloudera and Hortonworks merger completed. In June of 2019, IBM and Cloudera expanded partnership to include the entire Cloudera portfolio.
>
> CDP Private Cloud Base combines the best of Cloudera Distribution Hadoop (CDH) and Hortonworks Data Platform (HDP) functions and services.

Cloudera Data Platform Private Cloud Base

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

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

CDP Private Cloud Base consists of various components, such as Apache Spark, Apache Hive 3, and Apache HBase, along with many other components for specialized workloads. You can select any combination of these services to create clusters that address your business requirements and workloads. Several pre-configured packages of services are also available for common workloads.

With CDP Private Cloud Base supporting a separation of compute and storage design, integrating with IBM Spectrum® Scale provides the end-to-end solution to support the high demand workloads across different protocols. It also gives the ability to grow compute and storage requirements separately when doing analytics and AI in the same namespace.

IBM Spectrum Scale and Elastic Storage System

IBM Spectrum Scale is an industry-leading software for file and object storage. It can be deployed as a software-defined storage management solution that effectively meets the demands of AI, big data, analytics, and high-performance computing workloads. It has market leading performance and scalability, and a wealth of sophisticated data management capabilities.

IBM Elastic Storage System (ESS) is a fully integrated and tested Spectrum Scale storage building block that provides superb enterprise performance, reliability, availability, and serviceability. ESS is an optimum way to deploy Spectrum Scale storage for most Spectrum Scale use cases.

Integrated solution overview

CDP Private Cloud extends cloud-native speed, simplicity, and economics for the connected data lifecycle to the data center. It enables IT to respond to business needs faster and deliver rock-solid service levels so that users can be more productive with data.

CDP Private Cloud Base brings business value to enterprises by analyzing their disparate data sources and deriving actionable insights from them. This analytics journey typically starts with consolidation of different data silos to form an *Active Archive*. The Active Archive is then used to get a single view of the customer and perform further predictive analytics on them.

With IBM Spectrum Scale, clients can build highly scalable and globally distributed data lakes to form their Active Archives. IBM Spectrum Scale becomes the storage layer for your CDP Private Cloud Base environment as an alternative to native Hadoop Distributed File System (HDFS). It supports the access of the data by using HDFS Remote Procedure Calls (RPC) and is not apparent to the applications that use CDP Private Cloud Base. With IBM Spectrum Scale, you get more flexible deployment models for your storage system that help you optimize infrastructure costs.

IBM Spectrum Scale and CDP Private Cloud Base were first certified with IBM Spectrum Scale V5.1 and CDP 7. Since the first certification, CDP Private Cloud Base and IBM signed an agreement to certify both the products on an ongoing basis for their new releases and keep the certification current. (For more information about certified software levels, see Table 1 on page 27.) This certification is for IBM Spectrum Scale software and applies to all deployment models of IBM Spectrum Scale, including IBM Elastic Storage® System.

Benefits of integration

The following top benefits are realized by using IBM Spectrum Scale with CDP Private Cloud Base:

► Extreme scalability with parallel file system architecture

IBM Spectrum Scale is a parallel architecture. With a parallel architecture, no single metadata node can become a bottleneck. Every node in the cluster can serve both data and metadata, which enables a single IBM Spectrum Scale file system to store billions of files. This architecture enables clients to grow their CDP Private Cloud Base environments seamlessly as the data grows. Also, one of the key value propositions of IBM Spectrum Scale, especially with IBM Elastic Storage System (ESS), is running diverse and demanding workloads, plus the ability to tier down to Active Archive.

► A global namespace that can span multiple Hadoop clusters and geographical areas

Using IBM Spectrum Scale global namespace, clients can create active, remote data copies and enable real-time, global collaboration. This namespace enables global organizations to form *data lakes* across the globe, and host their distributed data under one namespace.

IBM Spectrum Scale also enables multiple Hadoop clusters to access a single file system while still providing all the required data isolation semantics.

The IBM Spectrum Scale Transparent Cloud Tiering feature can archive data into a S3/SWIFT compatible cloud object storage system, such as IBM Cloud® Object Storage, Microsoft Azure object storage service, or Amazon S3, by using the powerful IBM Spectrum Scale information lifecycle management (ILM) policies.

► A reduced data center footprint with the industry's best in-place analytics

IBM Spectrum Scale has the most comprehensive support for data access protocols. It supports data access by using NFS, SMB, POSIX, and HDFS. This feature eliminates the need to maintain separate copies of the same data for traditional applications and for analytics.

► True software-defined storage that is deployed as software or as a pre-integrated system

You can deploy IBM Spectrum Scale as software directly on commodity storage-rich servers or deploy it as part of a pre-integrated system by using the IBM Elastic Storage System to remote mount to the CES HDFS cluster. Clients can use software-only options to start small, while still using enterprise storage benefits. With IBM Elastic Storage System, clients can control cluster sprawl and grow storage independently of the compute infrastructure. IBM Elastic Storage System uses erasure coding to eliminate the need for the three-way replication for data protection that is required with other solutions.

► IBM hardware advantage

A key advantage for IBM Elastic Storage System is to lower capacity requirements. IBM Elastic Storage System requires 30% extra capacity to offer similar data protection benefits. IBM Power Systems servers along with the IBM Elastic Storage System offer the most optimized hardware stack for running analytics workloads. Clients can enjoy up to three times reduction of storage and compute infrastructure by moving to IBM Elastic Storage System compared to commodity scale-out x86 systems.

To support the security and regulatory compliance requirements of organizations, IBM Spectrum Scale offers Federal Information Processing Standard (FIPS) compliant data encryption for secure data at rest, policy-based tiering/ILM, cold data compression, disaster recovery, snapshots, and backup and secure erase. The CDP Private Cloud Base Atlas and Ranger components provide more data governance capabilities and the ability to define and manage security policies.

Benefits of separation of compute and storage

Deploy compute cluster and storage cluster separately is becoming popular primarily because it disaggregates storage from compute in Hadoop environment, which enables compute and storage to grow independently as per business requirements. This architecture significantly helps controlling the cluster sprawl and data center footprint.

Most of the commercial shared storage offerings require accessing the same data by using multiple access protocols on different file systems. Industry standard protocols access (for example, Windows SMB and NFS) enables organizations to build a common data lake for Hadoop and non-Hadoop applications. The adoption of containerized workloads is another reason why shared storage deployments are being considered.

A dedicated storage system (for example, ESS) can also provide powerful security enhancement, such as native encryption. It also can provide enterprise-level data management features, such as snapshot, compression, and disaster recovery. These features can be enabled in the storage system without affecting the compute cluster.

Component relationship

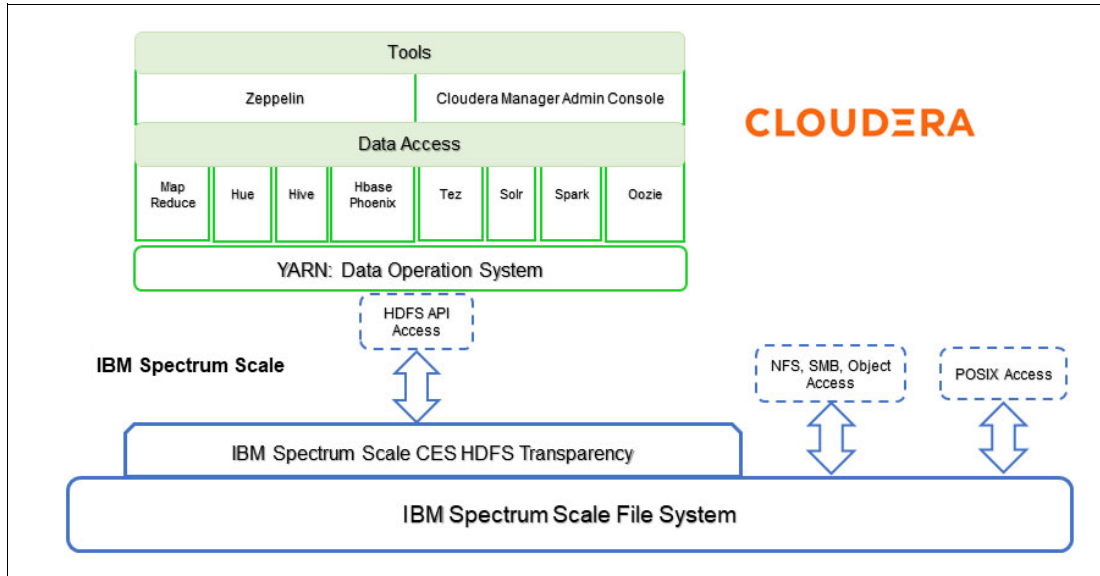Figure 1 shows the relationships between IBM Spectrum Scale and the CDP Private Cloud Base components.

Figure 1   CDP Private Cloud Base and IBM Spectrum Scale component relationship

Integration with Cloudera Data Platform Private Cloud Base

CDP Private Cloud Base (CDP PCB) consists of Cloudera Manager (CM) and Cloudera Data Hub (CDH) runtime components. CDP Private Cloud Base CM is certified starting at 7.2.3 and CDH at 7.1.4 on IBM Spectrum Scale 5.1.0.1. The CM 7.2.3 is a version that is specific for IBM Spectrum Scale integration.

CDP Private Cloud Base uses IBM Spectrum Scale custom service descriptor (CSD) to integrate with IBM Spectrum Scale. The CSD is a file that describes a product for use with Cloudera Manager. Cloudera Manager can then support configuration, distribution, and monitoring of that product.

CDP Private Cloud Base connects to IBM Spectrum Scale through a floating IP address called Cluster Export Services (CES) IP address. The CES IP address is on IBM Spectrum Scale protocol nodes. The IBM Spectrum Scale CES HDFS uses the CES IP as the NameNode IP address in a NameNode HA environment. This CES IP is used as the `dfs.namenode.rpc-address.<clustername>.nn1` value for CDP services to connect to the CES HDFS cluster. Only the `*.nn1` is required to be configured in NameNode HA because the CES IP is moved to the standby NameNode during failover in a CES environment.

CES HDFS

Starting from HDFS Transparency version 3.1.1 and IBM Spectrum Scale version 5.0.4.2, HDFS Transparency is integrated with the IBM Spectrum Scale installation toolkit and Cluster Export Services (CES). The integration of CES with HDFS Transparency (see Figure 2) is called IBM Spectrum Scale CES HDFS Transparency, or CES HDFS.
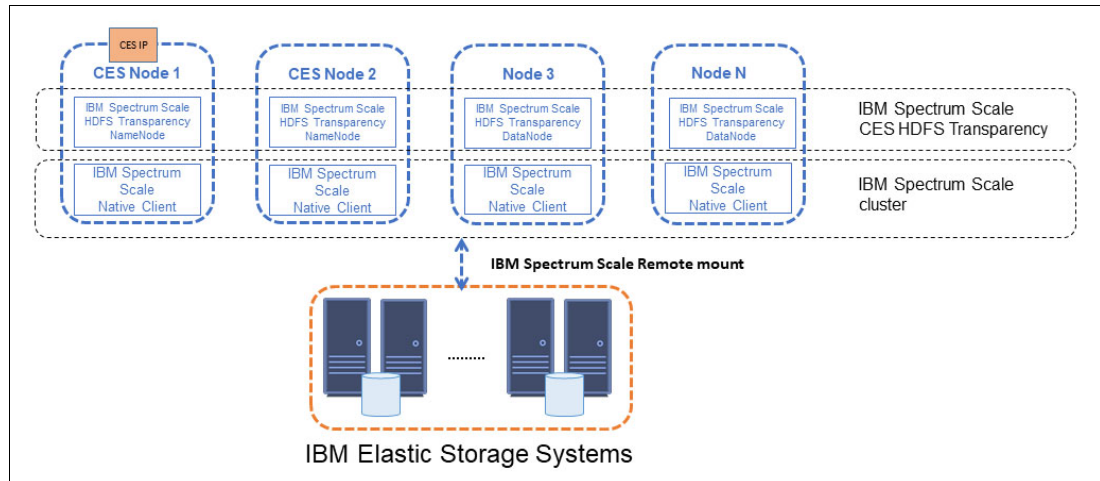

Figure 2   IBM Spectrum Scale CES HDFS Transparency cluster that uses ESS storage

IBM Spectrum Scale Cluster Export Services (CES) provides different protocol services, such as Network File System (NFS), Object, Hadoop Distributed File System (HDFS), or Server message Block (SMB) to an IBM Spectrum Scale cluster.

The CES infrastructure is responsible for the following tasks:

► Managing the setup for high-availability clustering that is used by the protocols. The participating nodes are designated as Cluster Export Services (CES) nodes or protocol nodes. The set of CES nodes is frequently referred to as the *CES cluster*. A set of IP addresses, the CES address pool (CES IP), is defined and distributed among the CES nodes. As nodes enter and leave the IBM Spectrum Scale cluster, the addresses in the pool can be redistributed among the CES nodes to provide high availability.

► Monitoring the health of these protocols on the protocol nodes and raising events or alerts during failures.

► Managing the addresses that are used for accessing these protocols by including failover and failback of these addresses because of protocol node failures. It is possible to use one IP address for all CES services. However, clients that use the SMB, Object, NFS, and Block protocols must not share the IP address for these protocols with the IP address that are used by the HDFS service to avoid affecting the clients of other protocols during an HDFS failover.

Only HDFS Transparency NameNodes are part of CES protocol nodes.

With the integration of HDFS into CES protocol, the use of the protocol server function requires extra licenses that need to be accepted.

The installation toolkit can install HDFS Transparency as part of the CES protocol stack. The CES interface can now control and configure HDFS Transparency by using the same interfaces as with the other protocols. The CES protocol manages the HDFS Transparency NameNodes only. The HDFS Transparency DataNodes are not part of the CES protocol nodes.

CES HDFS NameNode failover does not use ZKFailoverController. CES elects a new node to host the CES IP by using its own failover mechanism. HDFS clients always communicate with the same CES IP. Therefore, NameNode failover happens transparently. The Hadoop clients are required to be configured so that they know only about one IP address for the NameNode to connect to, even though a pool of NameNodes can exist in the CES HDFS cluster.

CES HDFS enabled with other protocol services recommendations

The recommended CES HDFS Transparency configuration is to remote mount to the ESS. To add other CES protocol to this environment, it is recommended to add them to the ESS as shown in Figure 3.
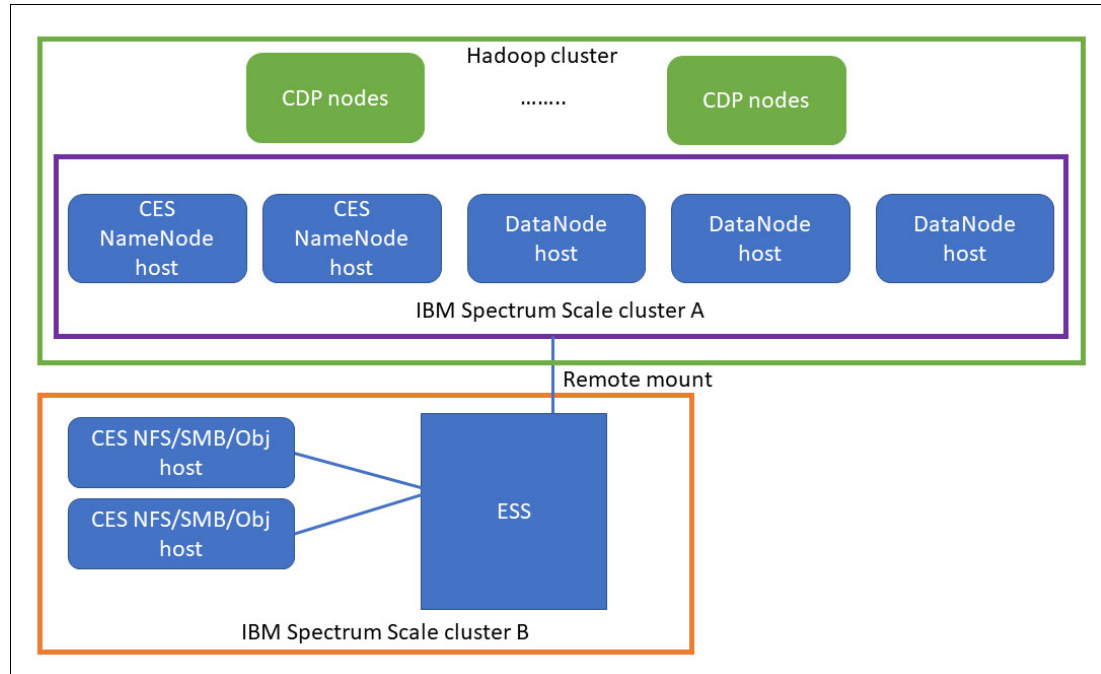

Figure 3   Recommended protocol configuration layout to the ESS

Consider the following recommendations and restrictions when CES HDFS protocol services are enabled along with other protocol node CES services, such as NFS, SMB, or Object:

► It is recommended to enable IBM Spectrum Scale CES HDFS in a dedicated IBM Spectrum Scale cluster and remote mount the file system from IBM ESS and separate the other CES protocols from HDFS. For more information about the recommended remote mount Spectrum Scale configuration, see "Deployment architecture" on page 15.

► All protocol nodes in the same IBM Spectrum Scale cluster must be on the same processor architecture. That is, all protocol nodes in a specific IBM Spectrum Scale cluster must be all x86, or all POWER Little Endian. IBM Spectrum Scale does not support an intermix of different processor architecture for protocol nodes in the same IBM Spectrum Scale cluster.

► Create specific CES IP addresses for CES HDFS usage.

► Each remote mount CES protocol node cluster must mount a dedicated CES shared root file system in the storage cluster. The CES share root file system cannot be shared between multiple remote mount protocol node clusters.

► The remote mount CES protocol node configuration does not support IBM Spectrum Scale Object or iSCSI services. For more information, see IBM Documentation.

► Review Hadoop ACL and IBM Spectrum Scale Protocols section in the IBM documentation for multi-protocol limitations.

IBM Spectrum Scale HDFS Transparency

IBM Spectrum Scale HDFS Transparency implementation integrates the NameNodes and the DataNodes services. It responds to the request as though it were HDFS on IBM Spectrum Scale file system.

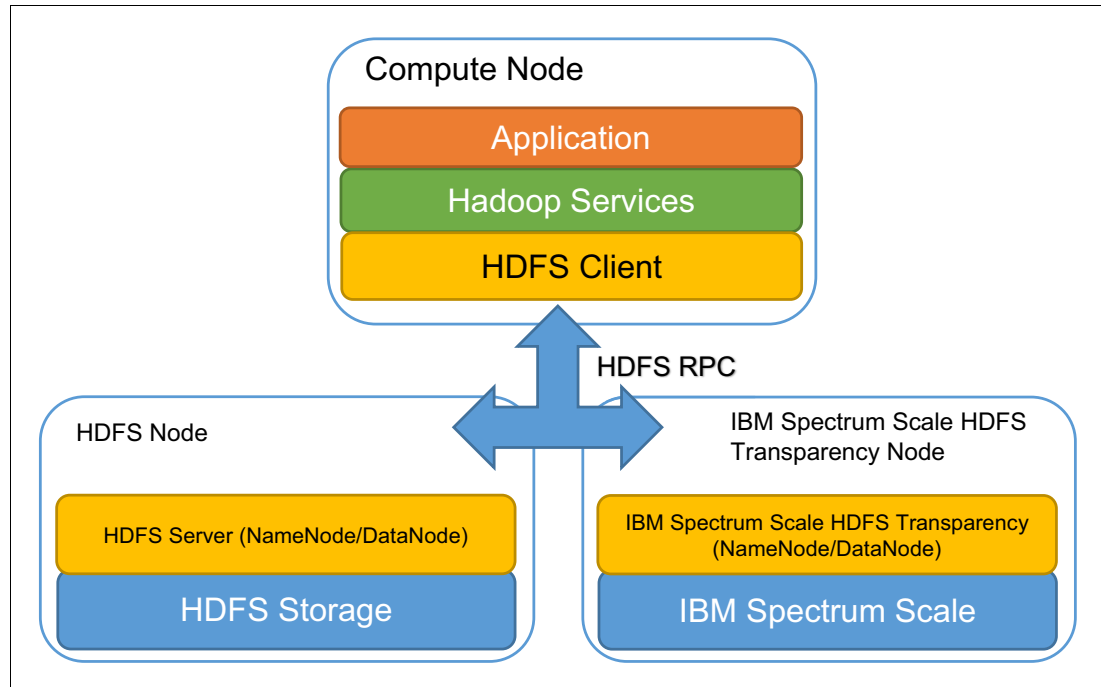Figure 4 shows the IBM Spectrum Scale HDFS Transparency component.

Figure 4   IBM Spectrum Scale HDFS Transparency

The use of HDFS Transparency includes the following advantages:

► Hadoop applications can run unmodified over IBM Spectrum Scale

► Immediate support for Hadoop applications and ISVs

► Single namespace for Hadoop and non-Hadoop workloads

► Reuse HDFS client as-is

► Stateless NameNode:

   – NameNode HA now uses edit log to store Kerberos credentials for long running jobs

   – The edit log resides in IBM Spectrum Scale so the information is accessible by the active and standby NameNode

► IBM Spectrum Scale is a distributed file system with distributed metadata and can scale to billions of files because of not having a centralized metadata NameNode as the bottleneck.

Cloudera Manager

Cloudera Manager (CM) is a Cloudera Hadoop administration tool with which users can manage, monitor, and configure multiple Hadoop clusters and its components by using the Cloudera Manager Admin Console web application or the Cloudera Manager API.

The Cloudera Manager Admin Console is a web application with which administrators and other Cloudera users can manage CDP Private Cloud Base deployments. By using the Cloudera Manager Admin Console, you can start and stop the cluster and other individual services, configure to add new services, manage security, and upgrade the cluster. You also can use the Cloudera Manager API to programmatically perform management tasks.

The heart of Cloudera Manager is Cloudera Manager Server. The Server hosts the Cloudera Manager Admin Console. The Cloudera Manager API and application logic is responsible for installing software, configuring, starting or stopping services, and managing the cluster that runs other Cloudera services. The Cloudera Manager components are shown in Figure 5.
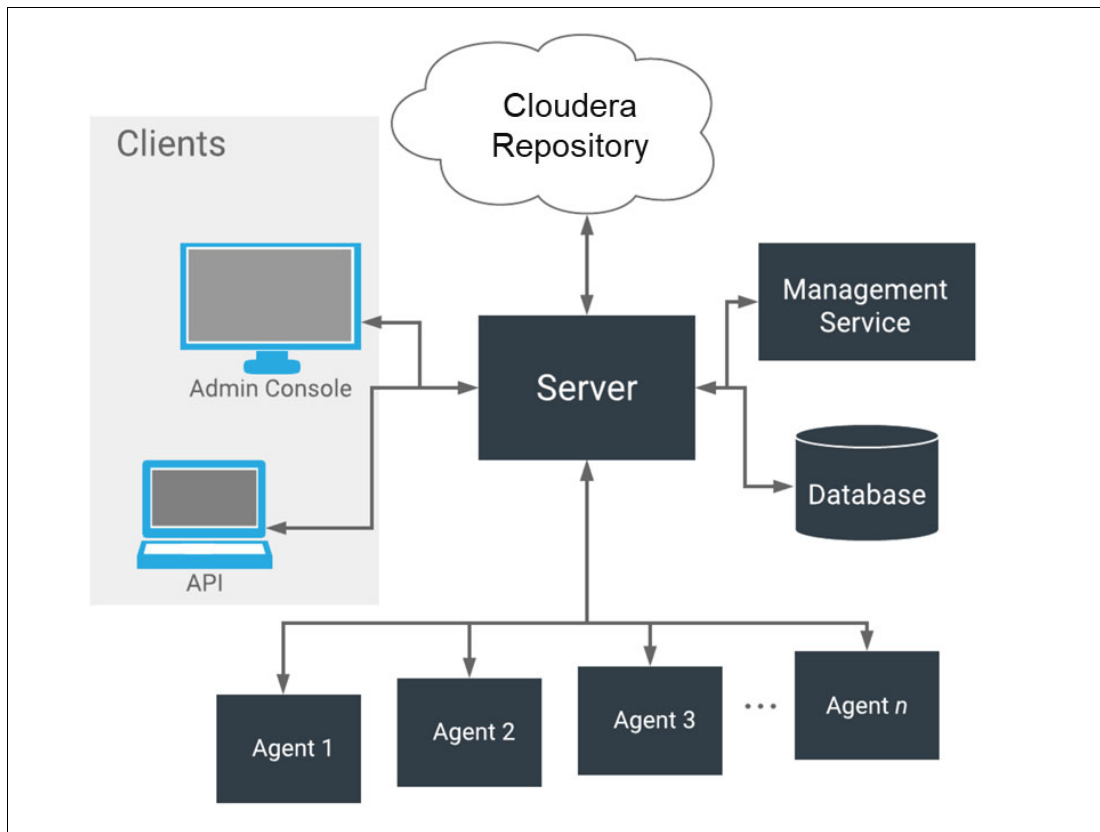

Figure 5   Cloudera Manager components

The Cloudera Manager server runs on a host in your CDP Private Cloud Base deployment. It manages the clusters by using Cloudera Manager Agents that run on each host in the cluster.

The CM Agent is a Cloudera Manager component that works with the Cloudera Manager Server to manage the processes that map to role instances.

Monitoring

When IBM Spectrum Scale HDFS CSD is integrated, the CM can start and stop the CES HDFS Transparency NameNodes and DataNodes, as shown in Figure 6.
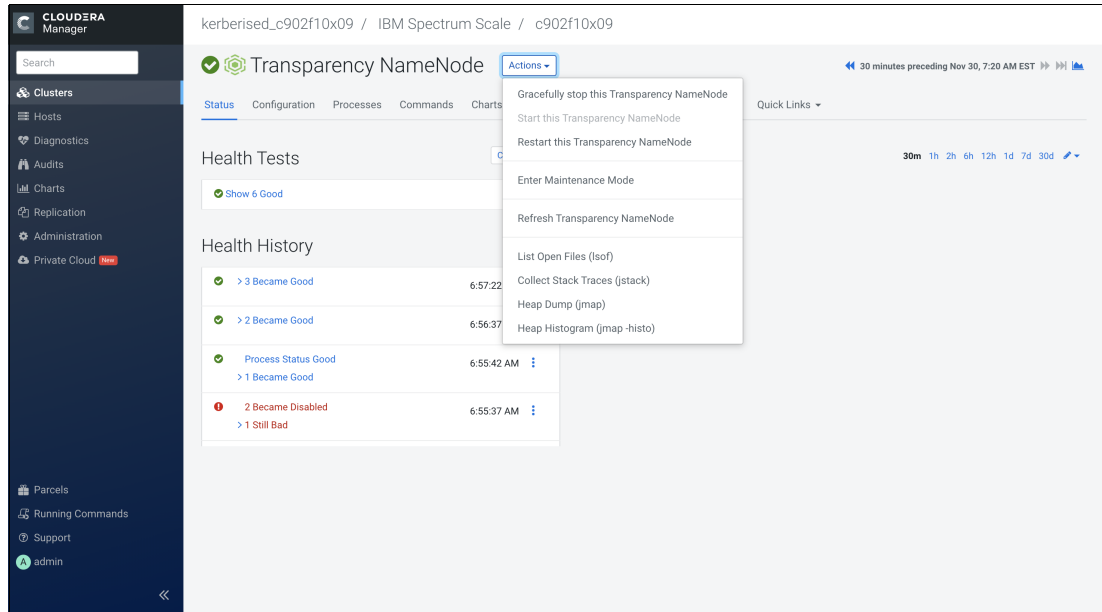

Figure 6   Using CM to start and stop the CES HDFS Transparency NameNodes

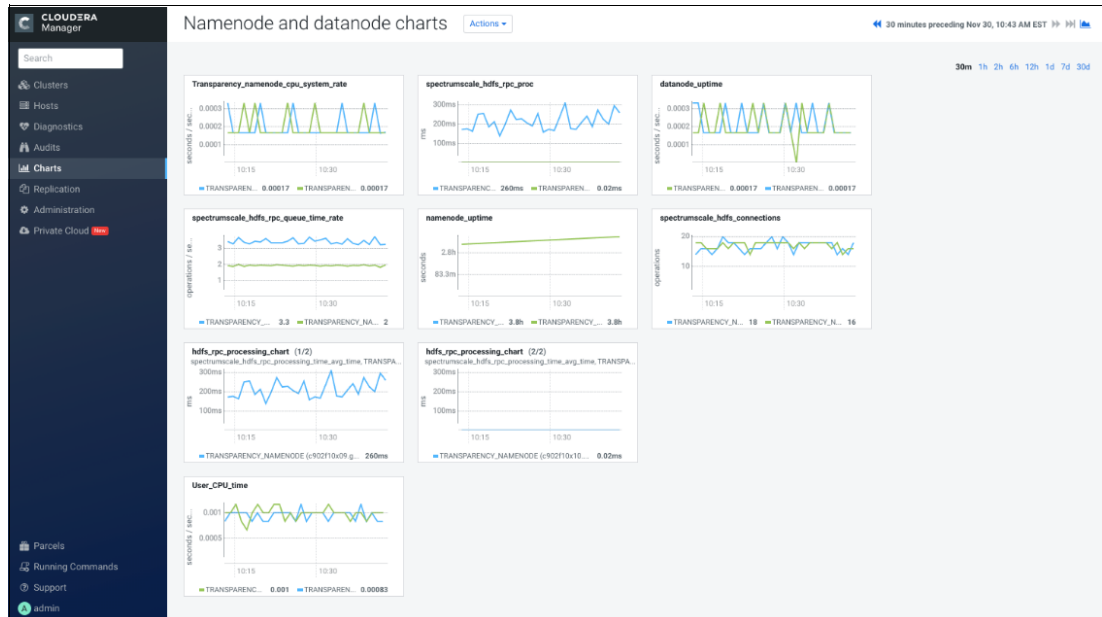The CM also shows HDFS Transparency metrics graph information (see Figure 7).

Figure 7   Using CM to show HDFS Transparency metrics graphs

Custom Service Descriptor

CM gives the ability to add your own managed service by using Custom Service Descriptors (CSDs). A third-party service that uses CSDs can use the features of Cloudera Manager, such as monitoring, resource management, configuration, distribution, and lifecycle management. This service appears in Cloudera Manager as any other service.

A CSD is linked to one service type in Cloudera Manager and is packaged and distributed as a `.jar` file. Cloudera Manager uses the CSD to know how to manage the deployed software (start and stop, configuration, resource management, and so on). A CSD is what provides the ability for a partner to have a service show up in the wizard and status pages. For more information, see this web page.

The IBM Spectrum Scale Custom Service Descriptors (CSD) integrates IBM Spectrum Scale HDFS transparency connector into the CM. This CSD file contains all of the configuration that is needed to describe and manage a Spectrum Scale service. The IBM Spectrum Scale CSD is provided in the form of a rpm installable that includes the IBM Spectrum Scale CSD `.jar` file embedded.

The IBM Spectrum Scale CSD package is provided by IBM and the `.jar` file is placed into `/opt/cloudera/cm/csd` where all CM CSD is stored.

The IBM Spectrum Scale CSD provides specific IBM Spectrum Scale parameters and commands so that the CM can manage, monitor, and connect to the CES HDFS Transparency cluster.

Security

This section describes various security and governance products that are supported by CDP Private Cloud Base and IBM Spectrum Scale.

Ranger and Atlas

Cloudera Runtime security and governance is managed by Apache Ranger and Apache Atlas:

► Apache Ranger manages auditing HDFS resources and access control through a user interface that ensures consistent policy administration in CDP clusters.

► Apache Atlas provides a set of metadata management and governance services that enable you to manage CDP cluster assets.

Ranger and Atlas are supported by IBM Spectrum Scale integration.

Kerberos

Kerberos is a network authentication protocol. It is designed to provide strong authentication for client/server applications by using secret-key cryptography. HDFS transparency supports full Kerberos and it is verified over CDP Private Cloud Base.

Data Encryption at rest

IBM Spectrum Scale offers built-in encryption support and provides support for file encryption that ensures both secure storage and secure deletion of data. IBM Spectrum Scale manages encryption through the use of encryption keys and encryption policies. Secure storage uses encryption to make data unreadable to anyone who does not possess the necessary encryption keys. The data is encrypted while at rest (on disk) and is decrypted on the way to the reader. Only data, not metadata, is encrypted.

It is important to understand the difference between HDFS encryption and built-in encryption with IBM Spectrum Scale. HDFS level encryption is per user-based whereas built-in encryption is per node-based.

HDFS encryption is supported starting in CDP Private Cloud Base 7.1.6 with IBM Spectrum Scale 5.1.1.

Transport Layer Security/Secure Sockets Layer encryption

Wire encryption protects data in motion, and Transport Layer Security (TLS) is the most widely used security protocol for wire encryption. TLS evolved from Secure Sockets Layer (SSL). TLS provides authentication, privacy and data integrity between applications communicating over a network by encrypting the packets transmitted between endpoints.

Transport Layer Security/Secure Sockets Layer (TLS/SSL) is supported starting in CDP Private Cloud Base 7.1.6 with IBM Spectrum Scale 5.1.1.

Multiple Hadoop clusters over the same file system

By using CES HDFS Transparency, you can configure multiple Hadoop clusters over the same IBM Spectrum Scale file system. For each Hadoop cluster, you need one CES HDFS Transparency cluster to provide the file system service.

As shown in Figure 8, an IBM Spectrum Scale file system services Hadoop cluster 1 and Hadoop cluster 2 at the same time through CES HDFS Transparency cluster 1 and CES HDFS Transparency cluster 2.
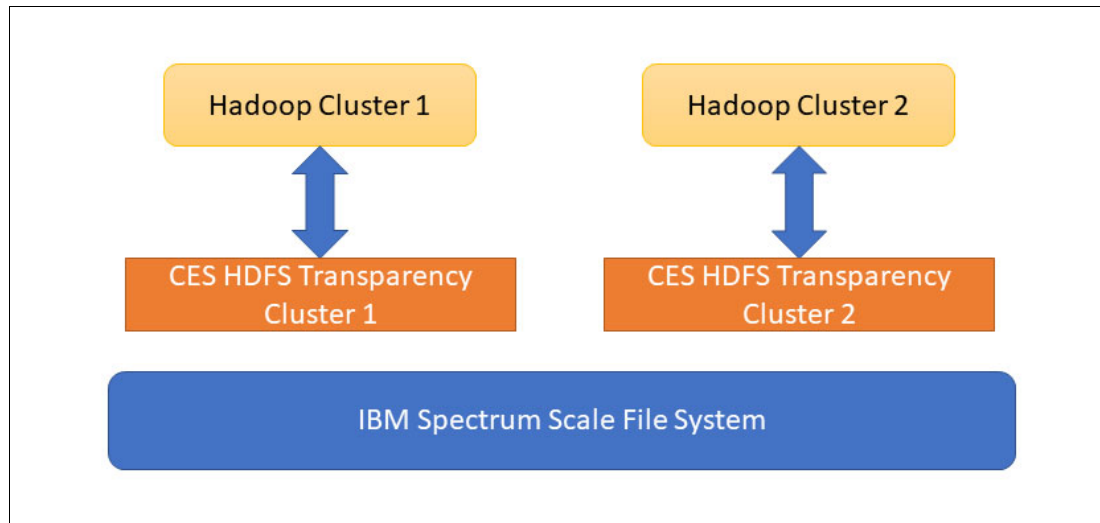

Figure 8   Two Hadoop Clusters over the same IBM Spectrum Scale file system

Consider the following key HDFS Transparency and IBM Spectrum Scale differences:

► If one file is set with Access Control List (ACL) (POSIX or NFSv4 ACL), IBM Spectrum Scale HDFS Transparency does not provide the interface to disable the ACL check at the IBM Spectrum Scale HDFS Transparency layer. If you want to disable the ACL for one file, the only way is to remove the ACL.

► HDFS level encryption is per user based, whereas IBM Spectrum Scale built-in encryption is per node based. Therefore, if the use case demands more fine-grained control at the user level, use HDFS level encryption. However, if you enable HDFS level encryption, you cannot get in-place analytics benefits, such as accessing the same data with HDFS and POSIX/NFS. This HDFS encryption is supported since HDFS Transparency 3.0.0-0 and 2.7.3-4.

► IBM Spectrum Scale provides its own caching mechanism that does not support HDFS caching. Caching that is done by IBM Spectrum Scale is more optimized and controlled,

especially when you run multiple workloads. The interface hdfs cache admin is not supported by IBM Spectrum Scale HDFS Transparency.

- ► NFS Gateway from native HDFS is not supported by IBM Spectrum Scale HDFS Transparency. IBM Spectrum Scale provides multiple protocol interfaces, including POSIX, NFS, and SMB. Customers can use IBM Spectrum Scale Protocol for NFS to access the data.

- ► The option distcp -diff is not supported for snapshot over IBM Spectrum Scale HDFS Transparency. Other options from distcp are supported.

- ► The interface from hdfs dfs is supported, whereas others (such as hdfs fsck) are not needed for IBM Spectrum Scale HDFS Transparency.

- ► HDFS file level COMPOSITE CRC file check sum is not supported.

For more information, see IBM Documentation.

Multiple IBM Spectrum Scale file systems support

Multiple IBM Spectrum Scale file system support is designed to give a single Hadoop cluster the ability to access two IBM Spectrum Scale file systems. It can access its own primary IBM Spectrum Scale file system and then can add in a secondary IBM Spectrum Scale file system to be accessed. The secondary file system can be from the same ESS or from a different ESS. The multiple IBM Spectrum Scale file system support helps resolve ViewFS support issues because ViewFS is not certified with Hive in the Hadoop community.

For more information, see IBM Documentation .

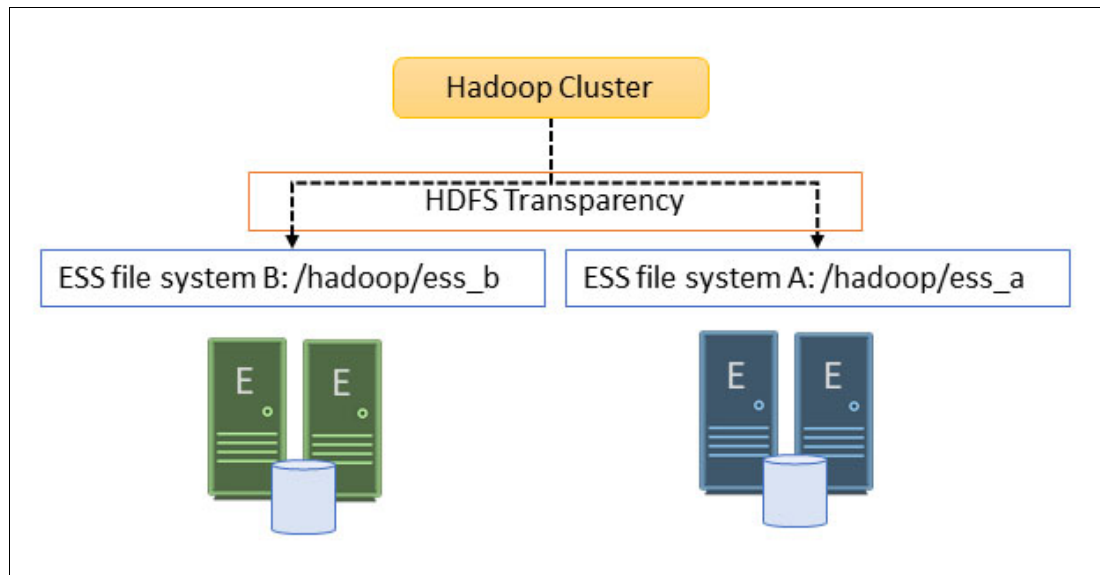As shown in Figure 9, an IBM Spectrum Scale ESS file system service the same Hadoop cluster through HDFS Transparency.

Figure 9   One Hadoop Cluster accessing multiple IBM Spectrum Scale file systems

Hadoop Storage Tiering with IBM Spectrum Scale

For customers that are adopting IBM Spectrum Scale and IBM Elastic Storage System (pre-integrated solution that is powered by IBM Spectrum Scale software) with Cloudera

Hadoop/Spark solution, a key requirement is to add IBM Elastic Storage System (ESS) into a Hadoop distribution cluster, such as CDH, HDP, and CDP. This feature eliminates the need to set up a separate Hadoop distribution cluster to gain the benefits of IBM ESS.

Hadoop Storage Tiering with IBM Spectrum Scale addresses this requirement. Enterprises that feature a standard CDP cluster with native HDFS can now add ESS as a storage tier in the same CDP cluster (see Figure 10). This configuration helps enterprises manage cluster sprawl by adding ESS-based shared storage to their CDP clusters.
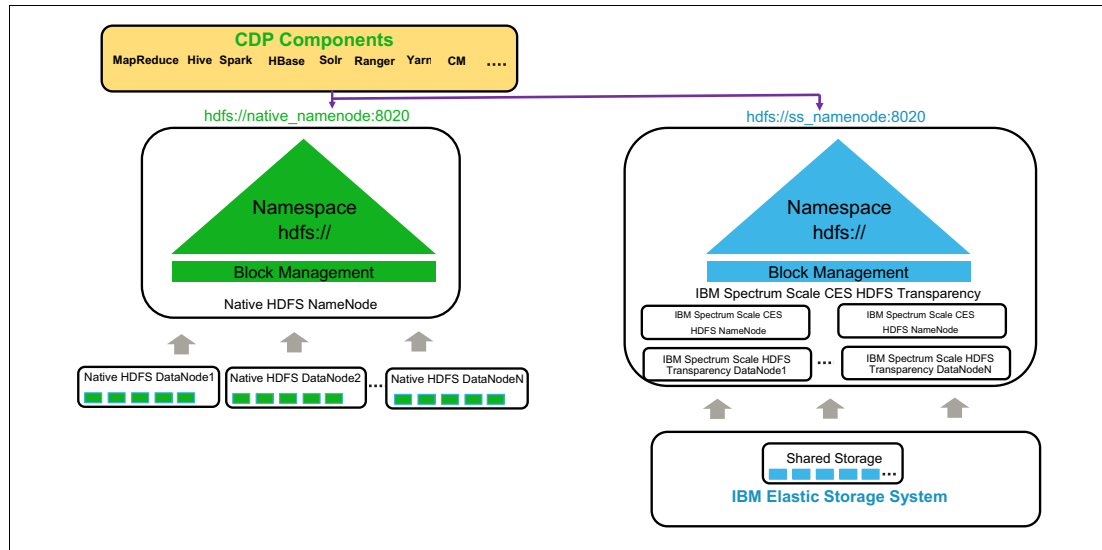

Figure 10   Hadoop storage tiering with IBM Spectrum Scale

This feature can be used in the following ways:

► As an ingest tier for faster ingest

   Enterprises can use IBM Spectrum Scale POSIX support with flash-based IBM ESS to get super-fast ingests for their Hadoop data lakes.

► As a secondary tier with shared storage

   Enterprises can use IBM ESS as a secondary tier in their Hadoop data lakes. This configuration enables them to grow storage independently of compute and also eliminates the need for three-way replication. The key benefit is the ability to run analytics directly on the secondary tier without having to bring the data into the primary tier.

► For data sharing between clusters

   If an enterprise wants to build a new analytics workflow on a new CDP cluster, but also needs access to the data from a CDP cluster, the tiering feature can enable this configuration without creating data copies. IBM ESS can be used as a secondary tier for the CDP cluster. The same ESS can act as the storage for a new CDP cluster. For example, some IBM customers are considering this scenario to introduce new IBM Power-based CDP clusters for demanding next generation analytics workflows.

13

► For migration to CDP Private Cloud Base with IBM Spectrum Scale

For enterprises that wants to migrate from their Hadoop cluster (as shown in Figure 11) to CDP Private Cloud Base with IBM Spectrum Scale, the tiering feature can help them move their data to the new environment. For Hadoop FPO environment, migration is the only path to move to CDP Private Cloud Base with IBM Spectrum Scale. Contact IBM or Cloudera professional services if you plan to migrate to CDP Private Cloud Base with IBM Spectrum Scale.
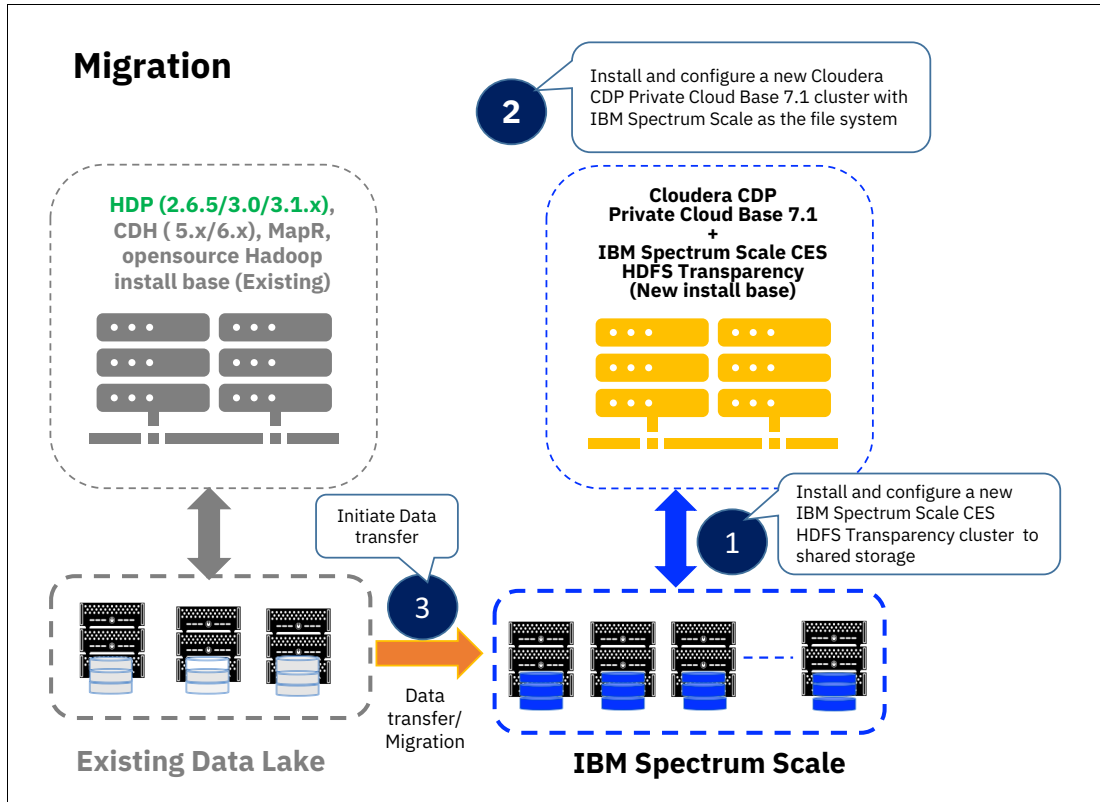

Figure 11   Migration to CDP Private Cloud Base with IBM Spectrum Scale using Hadoop storage tiering

Migration also gives customers a side-by-side migration path. By using this path, they can instantiate a new CDP Private Cloud Base with IBM Spectrum Scale cluster so that they can test their applications gradually. They also can move their workloads while retaining their current development and test or production cluster environment.

Figure 11 shows that customers with HDP, CDH, MapR, or open source Hadoop environments can move their data from the data lake to the new CDP Private Cloud Base with IBM Spectrum Scale file system with the help of an IBM or Cloudera professional service team.

For more information about Hadoop Storage Tiering with IBM Spectrum Scale, see IBM Documentation and this web page.

Deployment architecture

This section describes the deployment architecture of CDP Private Cloud Base with IBM Spectrum Scale integration system.

The recommended architecture is separation of Hadoop cluster hosts (master hosts, utility hosts, gateway hosts, or worker hosts) from Storage hosts (HDFS Transparency NameNodes and DataNodes). To achieve better performance, management, and enterprise-level storage capabilities, remote mount to the IBM ESS is the recommended deployment model for the CES HDFS cluster.

The benefits of separation of Hadoop cluster hosts from Storage hosts:

► Ability to manage Hadoop layer and Storage layer separately and by different teams

► Do not require to install IBM Spectrum Scale onto the Hadoop cluster hosts

► IBM Spectrum Scale requires specific kernel levels. The specific kernel levels are not required to be installed onto the Hadoop cluster nodes if the Hadoop cluster hosts are different than the IBM Spectrum Scale nodes (Storage hosts).

Shared Storage model

IBM Spectrum Scale allows Hadoop applications to access shared storage through the CES HDFS cluster. The shared storage can be IBM Elastic Storage System, Erasure Code Edition (ECE), or SAN-attached Shared Storage. Hadoop services or HDFS Transparency cannot be collocated with the ESS EMS, ESS IO nodes, or ECE nodes. IBM Elastic Storage System is a pre-integrated file storage solution that is powered by IBM Spectrum Scale software. This publication focuses on IBM Elastic Storage System-based deployments.

The CES HDFS cluster can be deployed by using the IBM Spectrum Scale Installation Toolkit. The installation toolkit supports a remote mount or single IBM Spectrum Scale file system model to the ESS.

Preferred IBM Spectrum Scale CES HDFS remote mount configuration

In the preferred remote mount configuration (as shown in Figure 12), the CES HDFS cluster is one IBM Spectrum Scale cluster and the ESS is another IBM Spectrum Scale cluster (which is often referred to as the *data cluster*).
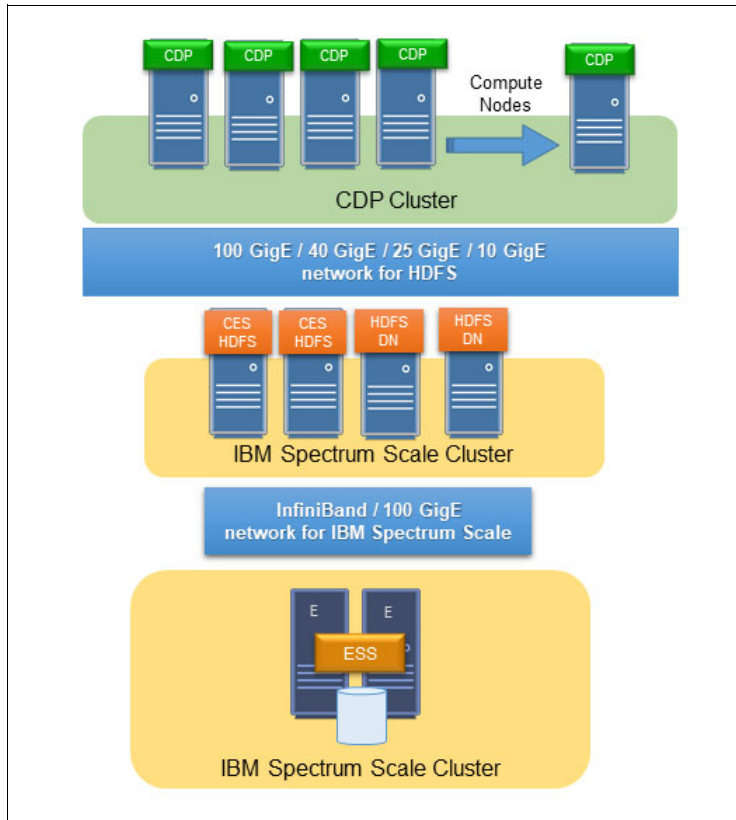
Figure 12   IBM Spectrum Scale Remote mount model

With this model, one IBM ESS data cluster can be shared with different groups and the remote mount configuration can isolate the data storage management from the IBM Spectrum Scale CES HDFS cluster. Therefore, stopping the IBM Spectrum Scale on a CES HDFS cluster does not stop the IBM Spectrum Scale on the ESS data cluster.

For full IBM Spectrum Scale GUI functions, the IBM Spectrum Scale GUI is set up on each of the IBM Spectrum Scale clusters separately for monitoring.

For more information about IBM Spectrum Scale remote mount configuration, usage, and considerations, see IBM Documentation.

Alternative IBM Spectrum Scale CES HDFS single cluster configuration

Circumstances exist in which CES HDFS protocol nodes might need to be implemented by using a single IBM Spectrum Scale cluster (as shown in Figure 13 on page 17). In this configuration, the CES HDFS nodes are deployed as the part of the same IBM Spectrum Scale cluster as the ESS.

Single cluster CES protocol nodes and ESS nodes in one IBM Spectrum Scale cluster is shown in Figure 13.
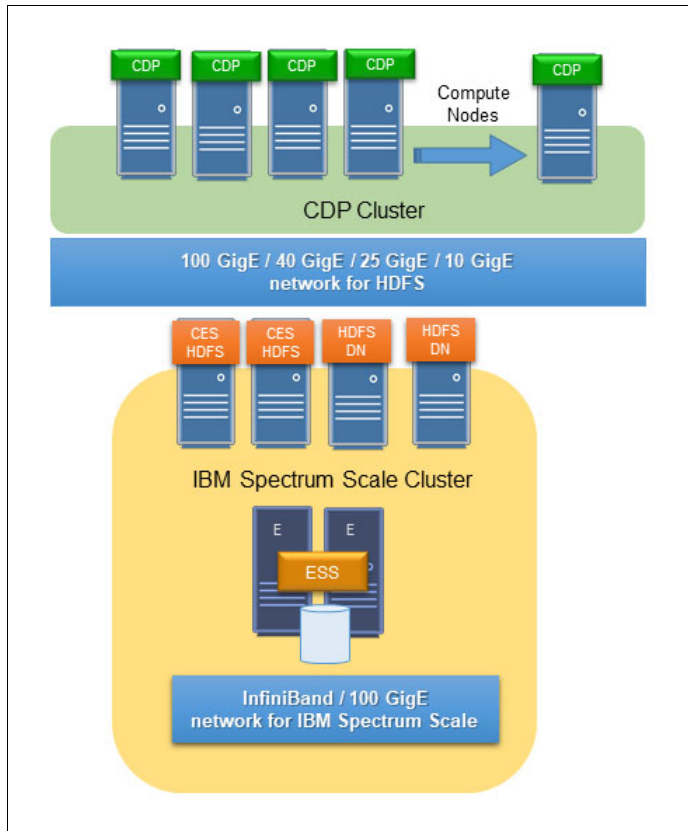
Figure 13   IBM Spectrum Scale as single storage system

The primary reason for deploying this single cluster Spectrum Scale configuration is if you want to manage and deploy all CES protocol nodes onto a single IBM Spectrum Scale cluster.

This viable single cluster CES HDFS configuration does mean you should plan for managing the combined IBM Spectrum Scale CES nodes and ESS data nodes as one cluster. Some operations and administration activities might affect all nodes in the one IBM Spectrum Scale cluster on both the CES HDFS nodes and the ESS nodes.

For more information about considerations and restrictions for choosing between these two configurations, see "Implementation guidelines" on page 17. Where possible, it is preferred to use the remote mount CES HDFS configuration that is shown in Figure 12 on page 16.

Implementation guidelines

The following sections describe architecture and implementation guidelines when CDP Private Cloud Base is implemented with IBM Elastic Storage System.

Cluster configuration

In a CDP Private Cloud Base and IBM Elastic Storage System deployment model, IBM Elastic Storage System serves as central back-end storage with a set of IBM Spectrum Scale CES HDFS Transparency nodes.

The CDP Private Cloud Base and IBM Spectrum Scale CES HDFS Transparency cluster are composed of CDP Nodes, CES protocol nodes (HDFS Transparency NameNodes), and IBM Spectrum Scale client nodes (HDFS Transparency DataNodes).

The CDP Node(s) depicted in the diagram are the Master, Utility, Gateway, and Worker Hosts. Because CES HDFS Transparency replaces the NameNode(s) and DataNodes, when looking at Cloudera Runtime Cluster Hosts and Role Assignments documentation, do not consider the NameNode under Master Hosts and the DataNode under the Worker Hosts columns (unless you are collocating the DataNode with other Hadoop services).

The configuration consists of IBM Spectrum Scale HDFS Transparency NameNodes and IBM Spectrum Scale Transparency DataNodes that are network-connected to the IBM Elastic Storage Systems. The recommended configuration is depicted in Figure 15 on page 19 with remote mount setup and Figure 14 shows the single cluster configuration.
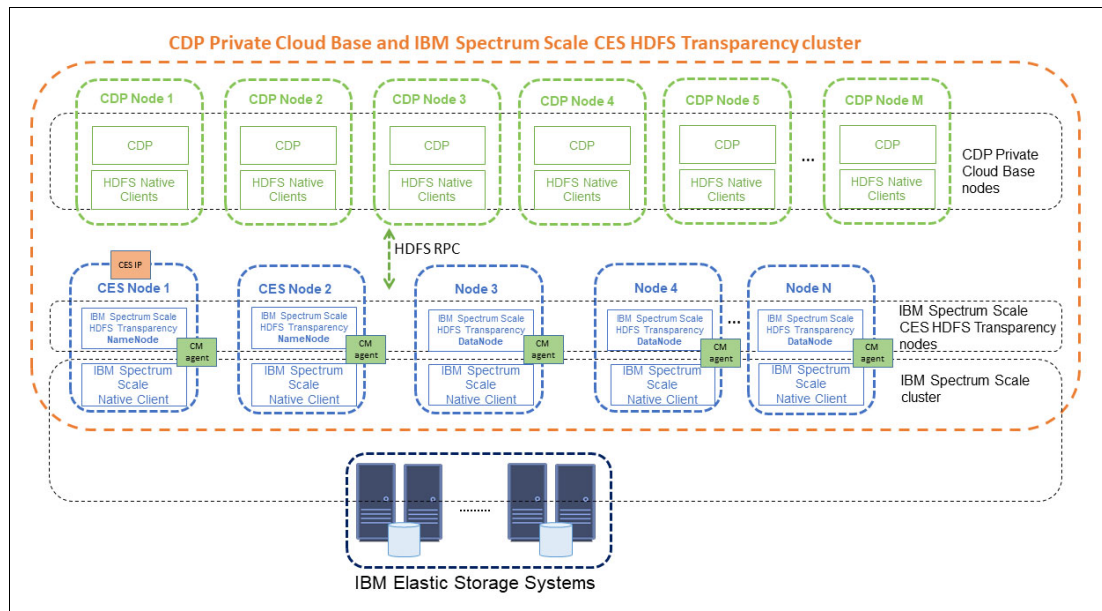

Figure 14   CDP Private Cloud Base and IBM ESS with single cluster configuration

IBM Spectrum Scale HDFS Transparency Name Nodes are managed by IBM Spectrum Scale CES protocol. IBM Spectrum Scale native client and Cloudera Manager (CM) agent are installed in all the IBM Spectrum Scale HDFS Transparency nodes. These nodes are the storage nodes.

The top section of Figure 14 indicates where the CDP Private Cloud Base services and HDFS native clients are installed. These nodes are the compute nodes that are separated from the storage nodes. The CDP Private Cloud Base nodes use HDFS RPC to access the IBM ESS through IBM Spectrum Scale HDFS Transparency layer.

Figure 14 also shows that the CES HDFS cluster and IBM ESS are part of the same IBM Spectrum Scale cluster.

One of the reasons for deploying the single cluster IBM Spectrum Scale configuration as shown in Figure 14 is if you need to concurrently deploy other CES protocol node services (such as Object or iSCSI) that are not supported by CES protocol nodes running in the remote mount configuration, as shown in Figure 15 on page 19.

For more information about restrictions for CES Protocol Node remote mount configuration, see IBM Documentation.
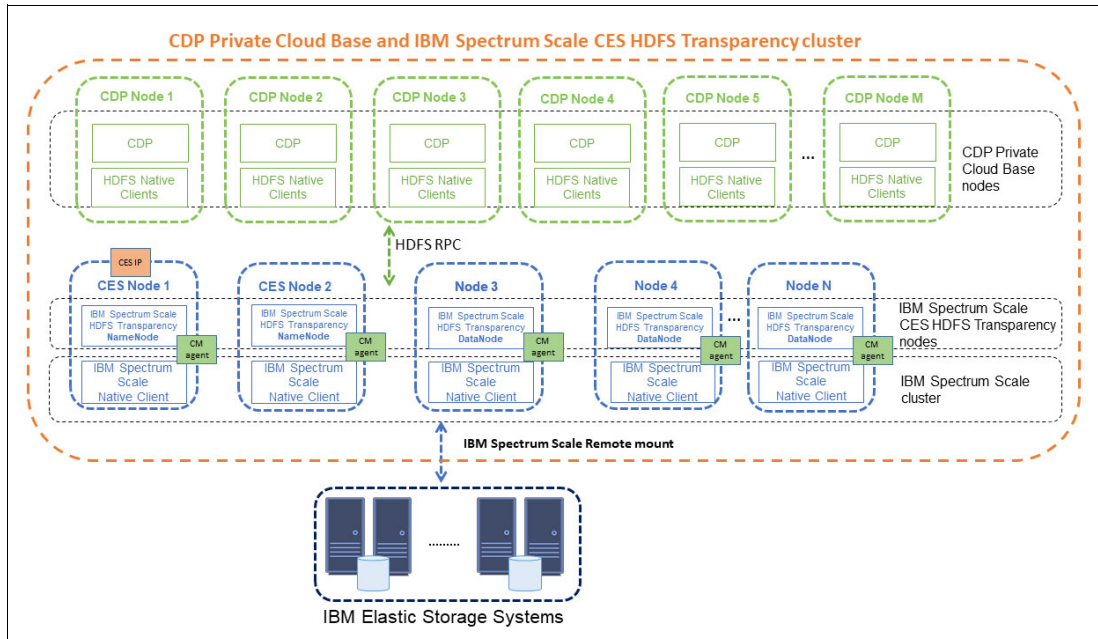
Figure 15   CDP Private Cloud Base and IBM ESS with remote mount cluster configuration

Figure 15 shows the CES HDFS cluster and IBM ESS that are part of different IBM Spectrum Scale clusters when remote mount mode is configured. When remote mount is used, an extra network is required to be set up between the CES HDFS cluster and the IBM ESS. Therefore, the clients access the IBM ESS through the network to the CES HDFS nodes and the CES HDFS nodes access to the IBM ESS data through another network.

The recommendation is to use the remote mount setup. Because multiple Hadoop clusters can use the same IBM ESS, the use of remote mount helps separate the IBM Elastic Storage System nodes and the different Hadoop clusters for better manageability.

For more information about remote mount, see IBM Documentation.

> **Note:** You must choose one of the configurations that are shown in Figure 14 on page 18 and Figure 15 for your CES HDFS environment, based on your overall requirements. The configurations are mutually exclusive of each other:
>
> ► If you have CES protocol nodes in a remote cluster, you cannot also have CES protocol nodes in the IBM Spectrum Scale ESS data cluster.
>
> ► If you have CES protocol nodes in the IBM Spectrum Scale ESS data cluster, you cannot also have CES protocol nodes in a remote IBM Spectrum Scale cluster.

Alternative cluster configuration

The recommended architecture is to separate the CDP Private Cloud Base nodes from the CES HDFS Transparency nodes using remote mount configuration to the storage as depicted in Figure 15 on page 19.

The following configurations are also supported:

► Collocate Hadoop services with HDFS Transparency DataNode. Note that HDFS Transparency NameNode still cannot collocate with other Hadoop services.

► Non-HA CES HDFS Transparency NameNode. This should only be used for proof of concept (POC) or non-production environments.

DataNode collocation configuration

HDFS Transparency DataNode can have other Hadoop services collocate within the same node.

Cloudera recommends installing specific services on the DataNode. Follow the worker hosts assignments column in the Cloudera Runtime Cluster Hosts and Role Assignments documentation for more information.

Collocation limitations:

► Cannot manage the Hadoop cluster hosts separate from the Storage hosts

► Requires to install IBM Spectrum Scale onto the Hadoop cluster hosts

► Requires specific kernel levels on the Hadoop cluster hosts

► IBM Spectrum Scale hosts require all uid/gid to be the same numeric value

► IBM Spectrum Scale requires passwordless ssh for either root or a non-root user with sudo privileges on all nodes

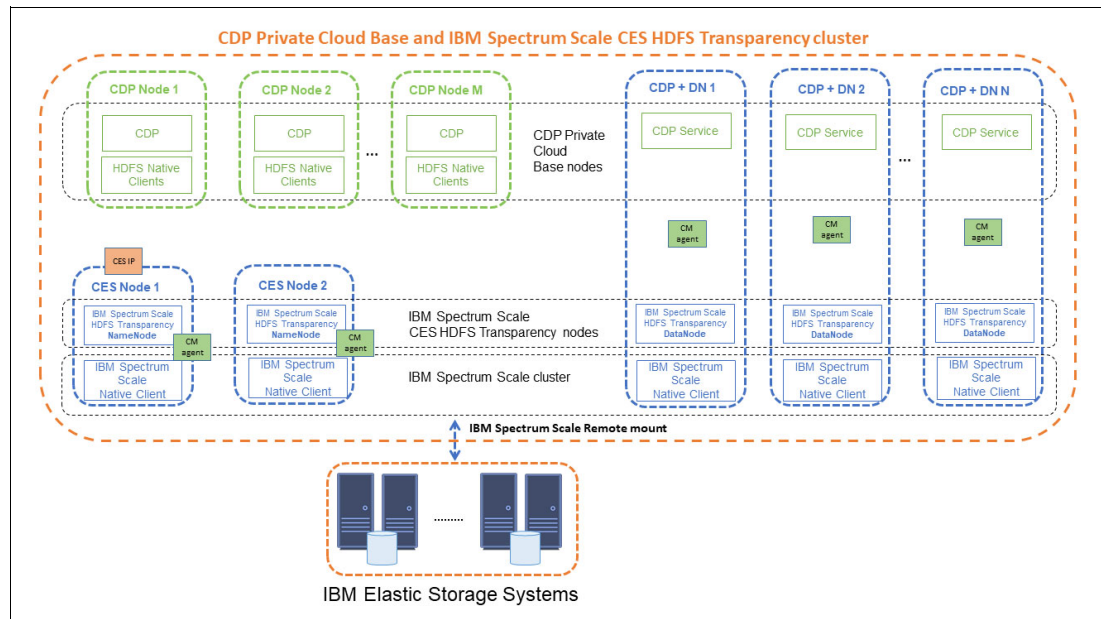Figure 16 shows an HA with DataNode collocation configuration.

Figure 16   HA with DataNode collocation configuration

Non-HA NameNode configuration

This configuration option should only be used for POC, dev, test, or non-production use cases.

If NameNode is not set up with high availability, then the Hadoop cluster will not be usable if the NameNode goes down.
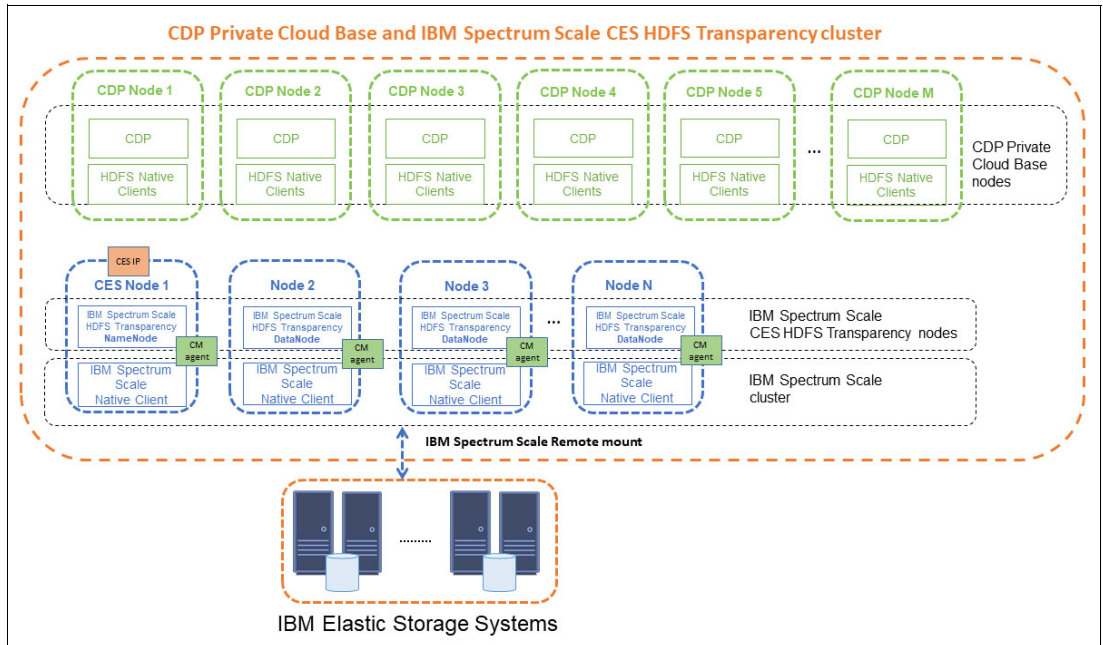
Figure 17 shows a non-HA configuration.

Figure 17   Non-HA configuration

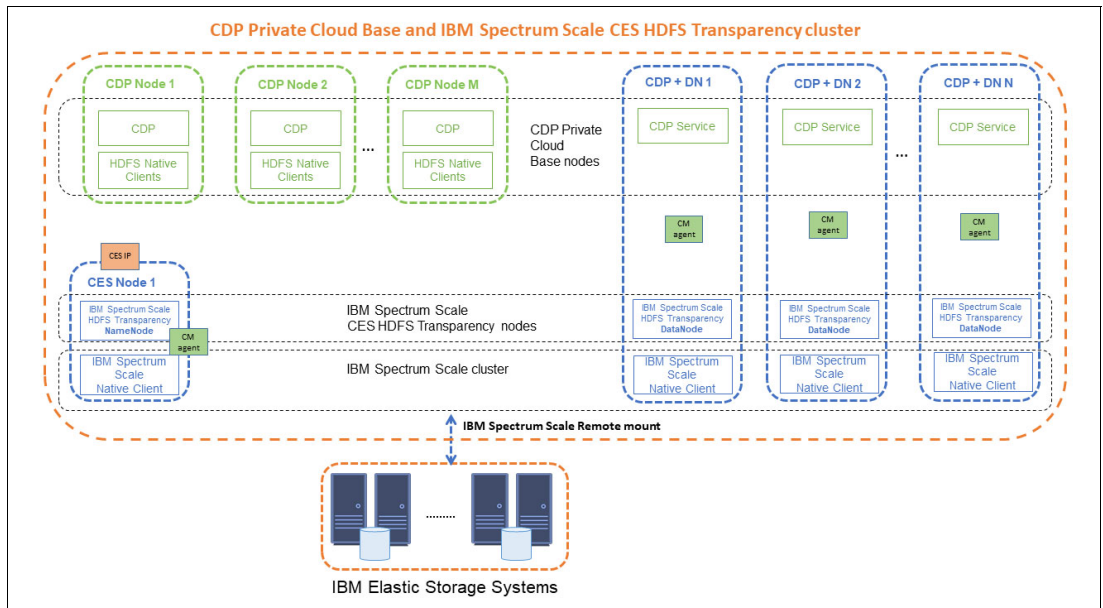Figure 18 shows a non-HA with DataNode collocation configuration.

Figure 18   Non-HA with DataNode collocation configuration

System design

In the architecture that is shown in Figure 19, the IBM Elastic Storage System is connected to a set of CES nodes. It is recommended to have two CES HDFS NameNodes for NameNode HA. Because resiliency and availability are important, CDP Public Cloud Base with IBM Spectrum Scale should include a NameNode HA setup for production.
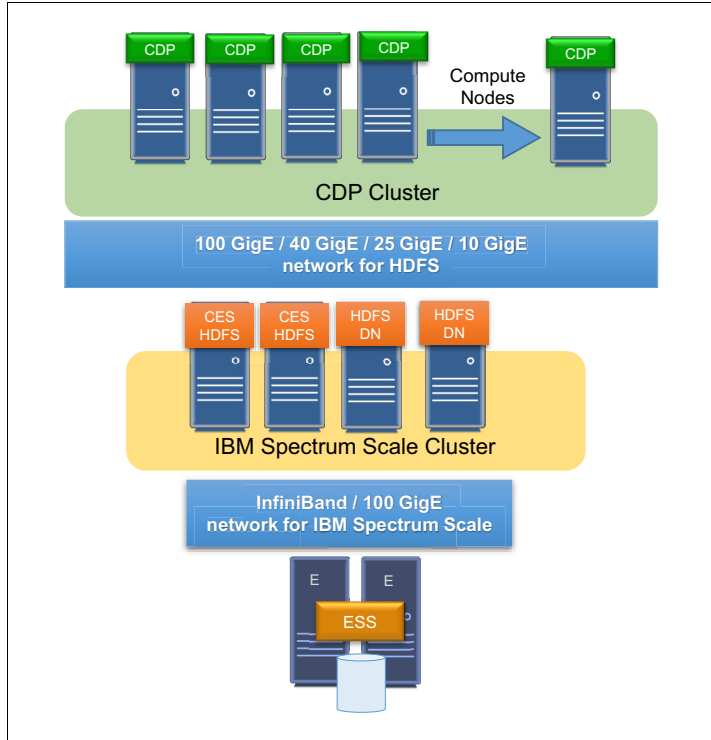

Figure 19   CDP Private Cloud Base and IBM Elastic Storage System with protocol nodes

If you plan to have other protocols in addition to HDFS, you must add CES nodes for their use (NFS and SMB).

In a separation of compute and storage architecture, the data flow for HDFS Transparency NameNodes and DataNodes are similar to native HDFS NameNodes and DataNodes to the Hadoop services and clients. The difference is when the Yarn Node Manager is not on the same node as the DataNode, the data flow features two network hops from the storage layer to the DataNode to the Yarn Manager node. The required data is sent over the network from the DataNode to the Yarn Manager node to be used for the job computation.

Figure 19 shows the IBM Spectrum Scale HDFS Transparency Name Nodes and a set of HDFS Transparency Data Nodes, through InfiniBand or 100 GigE network for IBM Spectrum Scale for better performance. A CDP Private Cloud Base cluster is connected to IBM Spectrum Scale HDFS Transparency through the network for HDFS, which supports 100 GigE, 40 GigE, 25 GigE, and 10 GigE. For more information about IBM Spectrum Scale configurations, see IBM Documentation.

These network and system design considerations exist regardless of whether you use the recommended remote mount configuration (see Figure 12 on page 16 and Figure 15 on page 19) or the single cluster configuration (see Figure 19 and in Figure 13 on page 17).

IBM Elastic Storage System models

IBM Elastic Storage System supports many high capacity and high IOPS model variations to fit your workload (see Figure 20). Select the model that best supports your overall capacity, performance, and availability requirements.
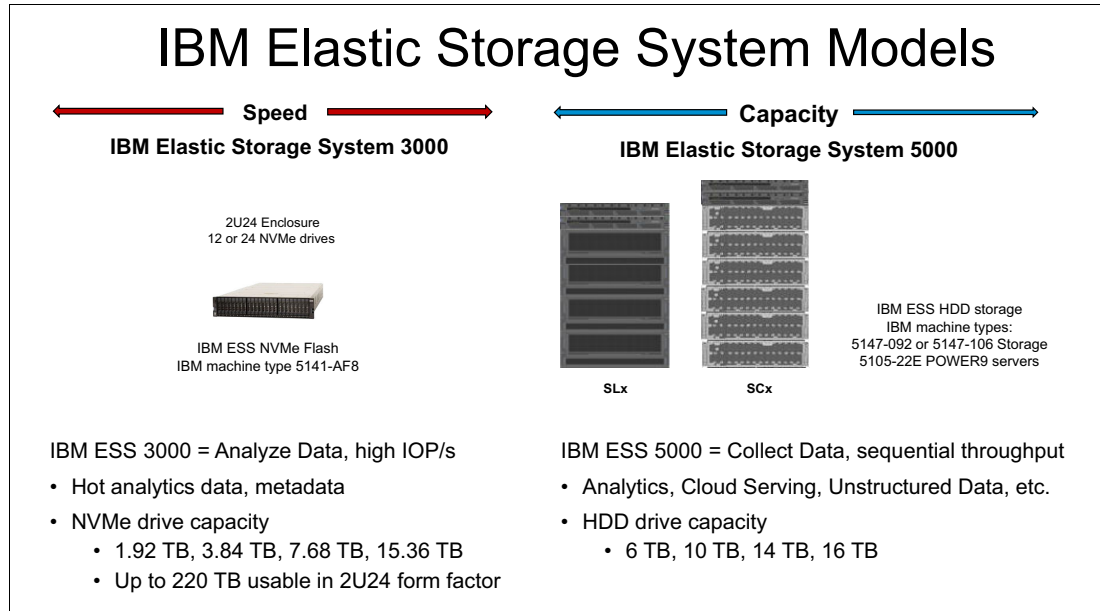

Figure 20   IBM Elastic Storage System models

Network

IBM Spectrum Scale consists of Admin and Daemon networks. If IBM Spectrum Scale Admin and Daemon networks are different, the specific network configurations for CDP Private Cloud Base, HDFS Transparency, and IBM Spectrum Scale cluster are recommended, as described next.

The `mmlscluster` command shows the Admin and Daemon node name information. The IBM Spectrum Scale Daemon node name and IP address fields correspond to the Daemon network that is used for data traffic in IBM Spectrum Scale and the Admin NodeName corresponds to the network that is used for running IBM Spectrum Scale administration commands (such as `mmlscluster` and `mmgetstate`).

In a dual network environment, two networks are used: Network 1 and Network 2. The following recommended network setup configuration options are available for the IBM Spectrum Scale cluster:

► Deploy Cloudera components, HDFS Transparency, CES IP/Hostname, and IBM Spectrum Scale Admin network in a common network; for example, Network 1.

The CDP Private Cloud Base service daemons (for example, *Yarn ResourceManager*) and HDFS Transparency daemons (for example, NameNode) should be in the same network to communicate with each other over RPC.

► Deploy IBM Spectrum Scale daemon network on the other network; for example, Network 2. Usually, this network is the high-speed network for IBM Spectrum Scale data traffic. Network 2 connects to the ESS.

For more information, see IBM Documentation.

IBM Elastic Storage System offers network adapter options. Each ESS data server two data servers in each ESS) provides three PCI slots that are reserved for high-speed data network adapters and one PCI slot is configured by default with a 4-port 1 GbE Ethernet adapter for management.

The three available high-speed network adapter slots are available to configure with any combination of Dual-Port 10/25 GigE, Dual-Port 100 GigE, or Dual-Port EDR InfiniBand adapters. Both ESS data servers must be configured with the same network adapter configuration.

For more information about updates to the 100 GigE or Enhanced Data Rate (EDR) InfiniBand adapters that are used in ESS, based on Mellanox ConnectX-5 network cards, see IBM Documentation.

Which high-speed network adapter you choose depends upon your performance requirements and networking infrastructure. In a 10/25 GigE network topology with IBM Elastic Storage System, carefully test the network bandwidth by using the free of charge, open source IBM Spectrum Scale network readiness tools to ensure that enough network bandwidth is available to meet performance expectations.

For all ESS models, a best practice is to use RDMA/EDR InfiniBand or 100 GigE high-speed topologies to interconnect the IBM Spectrum Scale/ESS storage data nodes and the CES HDFS protocol nodes. Otherwise, the performance benefits from an IBM Elastic Storage System building block likely are limited by the network connectivity between the IBM Elastic Storage System and the CES HDFS protocol nodes.

Note: It is important to have a reliable network for IBM Spectrum Scale to work optimally.

Data protection

IBM Elastic Storage System implements IBM Spectrum Scale erasure coding RAID software. IBM Spectrum Scale RAID implements sophisticated data placement and error-correction algorithms to deliver high levels of storage reliability, availability, and performance with cost-effective JBOD storage. For more information about IBM Spectrum Scale RAID and its components, see *IBM Spectrum Scale RAID Administration Guide*.

IBM Spectrum Scale RAID supports 2- and 3-fault-tolerant Reed-Solomon erasure codes and 2-, 3-, and 4-way replication. These configurations detect and correct up to one, two, or three concurrent faults, depending on the chosen RAID level.

Scaling

A primary advantage of the shared storage deployment model is its ability to grow storage performance and capacity independent of the compute infrastructure. If storage capacity or storage performance is insufficient, you can add storage into your cluster dynamically.

At the same time, you can add compute nodes without investing in capacity when the compute capacity is not sufficient. This granularity enables investment of resources based on your needs, as shown in Figure 21.
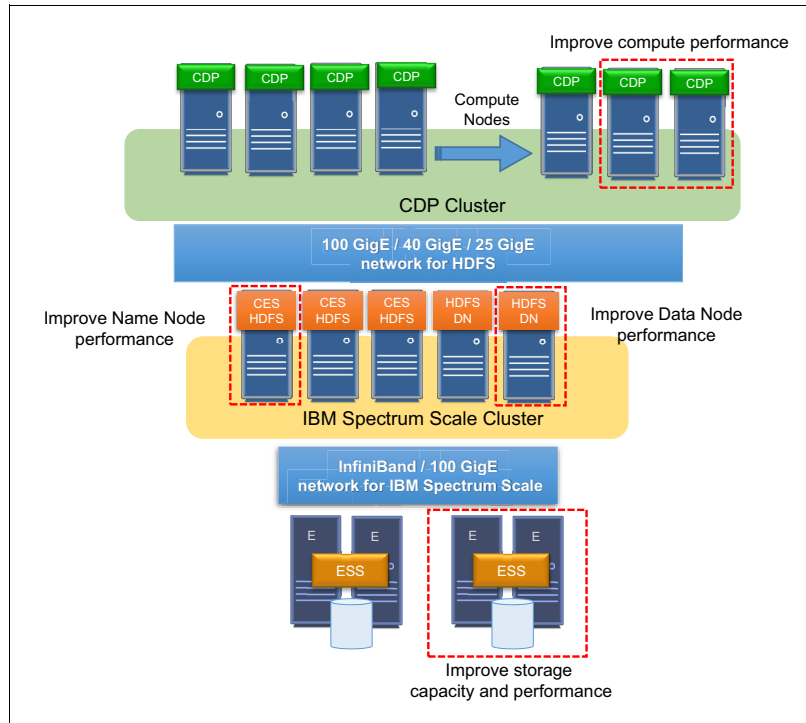

Figure 21   IBM Elastic Storage System scaling

If you want to have more storage capacity and performance, you can add another IBM ESS system. You also can add IBM Spectrum Scale HDFS Transparency Data Nodes into the cluster to improve bandwidth performance to the ESS. In the compute cluster, you can add CDP compute nodes to improve compute performance. All of these resources can be scaled out separately by your workload requirements.

Other preferred practices

Consider the following preferred practices while planning the deployment of CDP Private Cloud Base with IBM Elastic Storage Server:

▶ Tiering

IBM Spectrum Scale supports policy-based tiering and the ability to place metadata on separate storage from data. For performance-sensitive workloads, it is common to use solid-state storage for the file system metadata. For data, you can write policies to move file data to the flash tier for faster access. Policies can use many file attributes, including file heat, which enables you to create a policy based on how often the file is accessed, and not just on the last access.

For more information about IBM Spectrum Scale tiering, see IBM Documentation.

► File system block size

When creating a file system, design for two types of parameters: Parameters that can be changed after the file system is created and parameters that cannot. File system block size is the key parameter that must be determined at file system creation. After this parameter is set, the only way to change the block size is to re-create the file system.

In a IBM Spectrum Scale file system, you can store the file metadata (inode information) on the same storage as data or on separate storage. Consider the following options:

– Store file system metadata and data on separate storage. For more information, see IBM Documentation.

– The following preferred block sizes are used for Hadoop workloads on an IBM Elastic Storage System:

  • 1 MiB for a metadata only pool
  • 8 MiB for a data only pool

► IBM Spectrum Scale Hadoop performance tuning guide in IBM Documentation

System configuration

This section describes the minimum configuration setting when running CDP Private Cloud Base on IBM Spectrum Scale (see Figure 22).
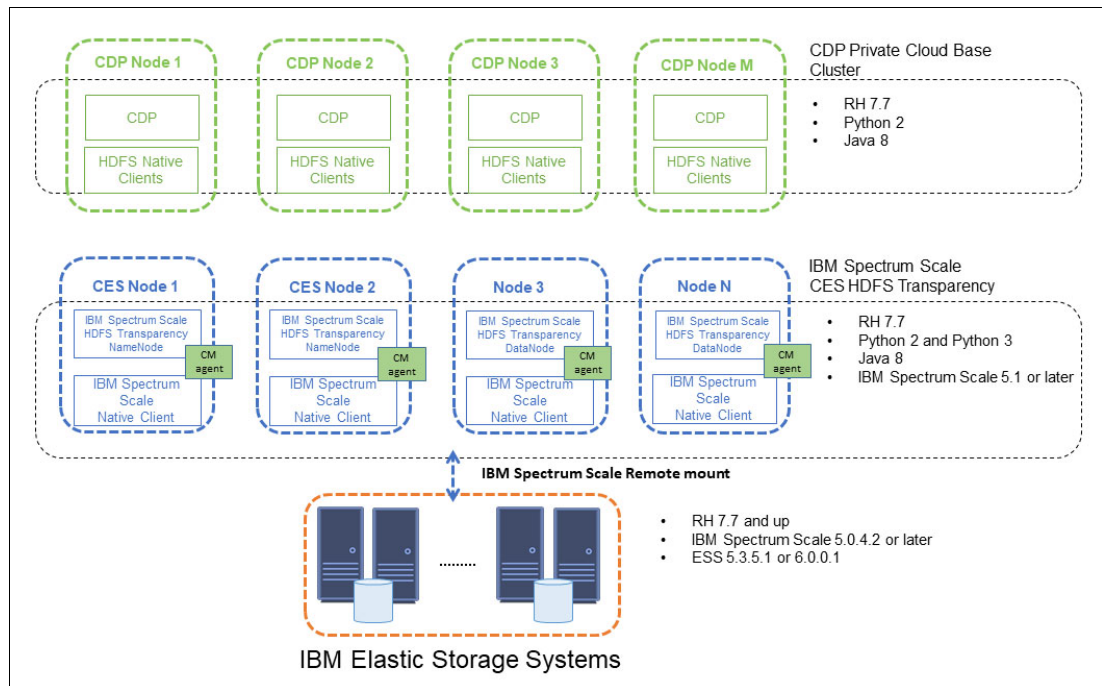

Figure 22   CDP Private Cloud Base on IBM Spectrum Scale minimum configuration

Minimum software version levels

Table 1 lists the minimum software version levels for CDP Private Cloud Base.

Table 1   Minimum software version levels for CDP Private Cloud Base

Component	Minimum release level	Description
Cloudera supported operating systems	64-bit Red Hat Enterprise Linux (RHEL) 7.7	► Supported operating system version for both CDP Private Cloud Base and IBM Spectrum Scale. ► CDP Private Cloud Base currently does not support RHEL 8.
Python	2.7	CDP Private Cloud Base currently does not support Python 3.
Java	Java 8/OpenJDK 1.8	► Supported Java version for CDP Private Cloud and IBM Spectrum Scale HDFS Transparency. ► HDFS Transparency currently does not support Java 11.
CDP Private Cloud Base	CM 7.2.3 CDH 7.1.4	IBM Documentation Big Data and Analytics support CDP Private Cloud Base Support Matrix for more information.

Table 2 shows the minimum software version levels for CES HDFS.

Table 2   Minimum software version levels for CES HDFS

Component	Minimum release level	Description
Operating system for IBM Spectrum Scale protocol and client	64-bit Red Hat Enterprise Linux (RHEL) 7.7	► Supported operating system version for both CDP Private Cloud Base and IBM Spectrum Scale. ► CM agent from CDP Private Cloud Base currently does not support RHEL 8.
Python	2.7 and 3	► Requires Python 2.7 and Python 3 to be installed. ► Python 2.7 is used for CM agent. ► Python 3 is used for IBM Spectrum Scale 5.1 or later.
Java	Java 8/OpenJDK 1.8	► Supported Java version for both CDP Private Cloud Base and IBM Spectrum Scale HDFS Transparency. ► HDFS Transparency currently does not support Java 11.
IBM Spectrum Scale HDFS CSD	1.0.0-0	The IBM Spectrum Scale Custom Service Descriptors (CSD) integrates IBM Spectrum Scale HDFS transparency connector into the CM.
IBM Spectrum Scale Big Data Analytics Integration Toolkit for HDFS Transparency (Toolkit for HDFS)	1.0.2.1	Used by IBM Spectrum Scale installation toolkit to deploy and install CES HDFS.
IBM Spectrum Scale CES HDFS Transparency Connector	3.1.1-3	IBM Documentation for IBM Spectrum Scale support for Hadoop CES HDFS

Component	Minimum release level	Description
IBM Spectrum Scale Client/HDFS DataNodes	5.1.0.1	IBM Documentation
IBM Spectrum Scale Protocol Nodes/CES HDFS NameNodes	5.1.0.1	IBM Documentation
IBM Spectrum Scale supported Linux and Kernel versions and hardware requirements for IBM Spectrum Scale Protocol services	N/A	IBM Spectrum Scale Frequently Asked Questions and Answers

Note: Consider the following points when setting up the CDP Private Cloud Base and IBM Spectrum Scale clusters:

► HDFS Transparency 3.1.1 is tightly coupled with IBM Spectrum Scale. The IBM Spectrum Scale Big Data Analytics Integration Toolkit for HDFS Transparency (Toolkit for HDFS), IBM Spectrum Scale HDFS Transparency, and IBM Spectrum Scale Cloudera Custom Service Descriptor (CDP CSD) are part of the IBM Spectrum Scale self-extracting installation package.

► CDP Private Cloud Base that is accessing HDFS Transparency nodes should use at least dual 10 Gb Ethernet or 25 Gb Ethernet connection.

► CES HDFS NameNode HA requires two or more CES nodes because a pool of NameNodes can be configured. Ensure to use a dedicated CES IP for HDFS protocol.

► For production, each NameNode x86 server should have a minimum of two sockets with at least eight cores each and at least 128 Gb of Memory for production clusters.

► For production, each DataNode x86 server should have a minimum of two sockets with at least eight cores each with 64 Gb of memory.

► For preferred performance, reserve 20% physical memory or up to 20 GB memory when you configure more than a 100-GB page pool for IBM Spectrum Scales.

► The protocol function (NFS/SMB) is software-only delivery; therefore, the capability and performance is based on the configuration that you choose. If you enable only one protocol, such as NFS, have a minimum of 1 CPU socket server with at least 64 GB of memory. If you enable multiple protocols or SMB, have a minimum two CPU socket server, with at least 128 GB of memory.

Table 3 lists the minimum software version levels for ESS.

Table 3   Minimum software version levels for ESS

Component	Minimum release level	Description
Protocol nodes	64-bit Red Hat Enterprise Linux (RHEL) 7.7	► Supported operating system version for CDP Private Cloud Base and IBM Spectrum Scale. ► CDP Private Cloud Base currently does not support RHEL 8. ► See Table 2 on page 27 for Minimum software versions for CES HDFS.
IBM Elastic Storage System	ESS 5.3.5.1 ESS 6.0.0.1	The following minimum ESS software levels support the minimum required IBM Spectrum Scale 5.0.4.2 level: ► IBM ESS 3000 specifications ► IBM ESS 5000 specifications ► Introducing IBM Spectrum Scale RAID ► IBM ESS 3000 and ESS 5000 I/O nodes and the ESS Management Server run Red Hat Enterprise Linux 8 with IBM Spectrum Scale 5

Support

CDP Private Cloud Base is certified with IBM Spectrum Scale starting at version 5.1.0.1 with CES HDFS Transparency starting with version 3.1.1-3.

Consider the following points regarding CDP Private Cloud Base integration with IBM Spectrum Scale:

► Requires an IBM Spectrum Scale CES HDFS cluster that is configured with shared storage before CDP Private Cloud Base is installed.

► A CES IP can be accessible from the CDP cluster.

► With separation of compute and storage, the CDP Private Cloud Base components are separated from the NameNodes and DataNodes, except for the CM agent.

► The CES HDFS cluster is recommended to have two NameNodes and at least three DataNodes for high availability.

► HDP 2.6 end of life was December 31, 2020 and HDP 3.1 end of life is December 31, 2021. For more information, see this web page.

► For more information about IBM Spectrum Scale, see IBM Spectrum Scale Planning Software requirements and FAQ.

► See CDP Private Cloud Base Runtime Cluster Hosts and Role Assignments for placement of Hadoop services onto the CDP Private Cloud Base hosts.

► CDP Private Cloud Base with IBM Spectrum Scale supports x86_64 and Power LE architectures. For more support matrix information, see IBM Documentation.

Limitations

Consider the following integration limitations:

► Short circuit reads should be disabled. The Node Manager is not on the same node as the DataNode.

► Do not share the IP address for other CES protocols with the IP addresses that are used by the HDFS service.

► Contact your account team for current upgrade information.

- The remote mount configuration is preferred for the CES HDFS Transparency cluster.
- Starting in IBM Spectrum Scale 5.1.0.1, IBM Spectrum Scale supports Object protocol from RHEL version 8.
- IBM Spectrum Scale Object protocol is not certified to be used through CDP Hadoop services. You can use the IBM Spectrum Scale object protocol through IBM Spectrum Scale nodes or via external services.
- Ensure that the CDP Private Cloud Base nodes and CES HDFS Transparency nodes are on the same operating system (OS) version and on the same architecture platform. Cloudera requires any node that installs its components to be on the same OS version and on the same architecture platform. The shared storage (for example, ESS) can be on a different OS and architecture platform. For example, if the CDP nodes and CES HDFS protocol nodes are on x86_64, the ESS can be on Power.

For more information about limitations, see IBM Documentation.

Additional references

- Cloudera:

  https://www.cloudera.com
- Cloudera and IBM:

  https://www.cloudera.com/partners/solutions/ibm.html
- IBM and Cloudera partnership:

  https://www.ibm.com/analytics/partners/cloudera
- Cloudera Blog: CDP Data Center: Better, Safer Data Analytics from the Edge to AI:

  https://blog.cloudera.com/cdp-data-center-better-safer-data-analytics-from-the-edge-to-ai
- Cloudera Docs - CDP Private Cloud Base (Private Cloud):

  https://docs.cloudera.com/cdp-private-cloud-base/latest/index.html
- Cloudera Docs - CDP Private Cloud Base (Private Cloud), Cloudera Manager:

  https://docs.cloudera.com/cdp-private-cloud-base/latest/concepts-cloudera-manager.html
- CSD Overview:

  https://github.com/cloudera/cm_ext/wiki/CSD-Overview
- Big data and analytics support:

  https://www.ibm.com/docs/en/spectrum-scale-bda?topic=big-data-analytics-support
- CES HDFS troubleshooting:

  https://www.ibm.com/docs/en/spectrum-scale-bda?topic=determination-ces-hdfs-troubleshooting
- IBM Spectrum Scale Hadoop performance tuning guide:

  https://www.ibm.com/docs/en/spectrum-scale-bda?topic=spectrum-scale-hadoop-performance-tuning-guide
- IBM Documentation for IBM Spectrum Scale:

  https://www.ibm.com/docs/en/spectrum-scale

- ► IBM Documentation for IBM Spectrum Scale FAQ:

  https://www.ibm.com/docs/en/STXKQY/gpfsclustersfaq.html

- ► IBM Documentation for IBM Elastic Storage System:

  https://www.ibm.com/docs/en/ess-p8

- ► *Implementation Guide for IBM Elastic Storage System 3000*, SG24-8443:

  http://www.redbooks.ibm.com/abstracts/sg248443.html

- ► *Introduction Guide to the IBM Elastic Storage System*, REDP-5253:

  http://www.redbooks.ibm.com/abstracts/redp5253.html

- ► *IBM Spectrum Scale Security*, REDP-5426:

  http://www.redbooks.ibm.com/abstracts/redp5426.html

- ► Workflow of a Hadoop Mapreduce job with HDFS Transparency & IBM Spectrum Scale:

  https://community.ibm.com/community/user/storage/blogs/chinmaya-mishra1/2020/11/23/workflow-of-a-mapreduce-job-with-hdfs-transparency

- ► I/O Workflow of Hadoop workloads with IBM Spectrum Scale and HDFS Transparency:

  https://community.ibm.com/community/user/storage/blogs/chinmaya-mishra1/2020/11/19/io-workflow-hadoop-hdfs-with-ibm-spectrum-scale

- ► IBM Documentation - IBM Spectrum Scale Protocol quick overview:

  https://www.ibm.com/docs/en/spectrum-scale/5.1.0?topic=quick-reference

- ► Kerberos: The Network Authentication Protocol:

  https://web.mit.edu/kerberos

Authors

This paper was produced by a team of specialists from around the world working at IBM Redbooks, Tucson Center.

Wei Gong is a Senior Software Engineer in IBM responsible for IBM Spectrum Scale development and client adoption. He has over 9 years of development on IBM Spectrum Scale core functions. Wei takes significant time with clients on IBM Spectrum Scale solution design, deployment, and performance turning. Wei has 5 years of storage development experience, including virtual machine storage system and storage HBA driver.

Linda Cham is a Senior Software Engineer in Poughkeepsie, NY. She has 4 years of experience in IBM Spectrum Scale Big Data Analytics (BDA) and 3 years in Life Science solutions and many years in High Performance Computing (HPC) development and service. Linda is the scrum master for the BDA worldwide teams and manages many projects through its lifecycle and customer engagements and advocacy. She is the author of multiple Big Data and Analytics blogs and videos, and is a consultant on BDA and Life Science topics.

Prashanth Shetty is a software Engineer in IBM working on testing and automation platform for IBM Spectrum Scale Big Data and Analytics products. He has 6 years of experience in IBM on various IBM Spectrum Scale product lines, such as IBM Spectrum Scale Integration with Hortonworks Data Platform, CES HDFS integration with IBM Spectrum Scale Installation toolkit, and deployment of IBM Spectrum Scale by using the installation toolkit. He holds a masters degree in Digital Communications from NITTE, Karnataka. Before coming to IBM, he worked for two years validating firmware and drivers for host bus adapters and MegaRAID controller cards for IBM and Dell Perc servers.

John Sing is Offering Evangelist for IBM Spectrum Scale, Elastic Storage Server. In his over 25 years with IBM, John has been a world-recognized IBM speaker, author, and strategist in enterprise storage, file + Object Storage, internet scale workloads and data center design, big data, cloud, IT strategy planning, high availability, business continuity, and Disaster Recovery. He has spoken at over 40 IBM conferences worldwide and is the author of eight IBM Redbooks® publications and nine IBM Redpaper publications.

Thanks to the following people for their contributions to this project:

Larry Coyne
IBM Redbooks, Tucson Center

Uday Kodoly
Chinmaya Mishra
Xin Wang
Piyush Chaudhary
Dave McDonnell
Bill Martinson
IBM Systems

Farzana Kader
David Fowler
Cloudera, Inc.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ► Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- ► Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- ► Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

35

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®	IBM Elastic Storage®	Redbooks®
IBM®	IBM Spectrum®	
IBM Cloud®	POWER®	

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

ibm.com/redbooks