Cloudera Public Cloud Reference Architecture

Date published: 2024-11-20
Date modified: 2024-11-20

https://docs.cloudera.com/

Legal Notice

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Abstract

The Cloudera Public Cloud disaster-recovery reference architecture is a high-level design and best practices guide to deploy Cloudera Public Cloud and to implement disaster-recovery use cases.

This document extends the disaster-recovery reference for CDP Private Cloud Base to cover Cloudera Public Cloud disaster-recovery use cases. The document discusses how you can leverage Cloudera tools and/or Cloud providers' native supportability to achieve your cloud-related disaster-recovery use cases for your workloads.

This document focusses primarily on the following two tiers:

- Tier 4: Point-in-time recovery - This tier is between one or more separate clusters in different regions. HDFS usually achieves this tier.
- Tier 5: Two-site commit/transition integrity - In this tier, data is continuously transmitted to primary and alternate backup sites. The cloud storage is often a fit for this tier.

To understand the tiers, see seven tiers of disaster recovery.

The document focuses on HDFS, Hive, and HBase data, metadata, access policies using Ranger, and lineage using Atlas.

Note: Cloudera Reference Architecture guides illustrate sample cluster configurations and certified partner products. The reference architecture guides are not replacements for official statements of supportability. They provide guidance to assist with deployment and sizing options. Statements regarding supported configurations in the reference architecture guides are informational and should be cross-referenced with the latest documentation.

Cloudera Public Cloud architecture might significantly affect the node and network sizing considerations. This reference architecture is appropriate for aggregated workload clusters running Cloudera Runtime.

Why use cloud disaster recovery?

Cloud disaster recovery provides organizations with the ability to recover their IT systems and data in the event of a disaster, such as hardware failure, natural disasters, cyberattacks, or human error, and also helps to fulfill compliance requirements.

The following use cases illustrate the need for cloud disaster recovery:

Data backup and recovery

Cloud disaster recovery allows organizations to regularly back up their critical data and applications to the cloud. In the event of data loss or corruption, they can quickly restore the data from the cloud backup.

Business continuity

Cloud disaster recovery ensures that the critical business processes can continue uninterrupted.

Protection from natural disasters

Organizations located in areas prone to natural disasters such as hurricanes, earthquakes, and floods can use cloud disaster recovery to replicate their data and applications to geographically distant cloud regions, ensuring data availability.

Ransomware and cybersecurity threats

Cyberattacks, such as ransomware, can encrypt or compromise data. Cloud disaster recovery enables organizations to recover their data from uninfected backups in the cloud, reducing the impact of such attacks.

Testing and DevOps

Cloud disaster recovery environments can be used for testing and development purposes without impacting the production environment. This ensures that the recovery processes are regularly tested and optimized.

Migrating to the cloud

When organizations are in the process of migrating their workloads to the cloud, they can use cloud-based disaster recovery to protect their data during the migration process. This minimizes downtime and risk.

Regulatory compliance

Many industries have stringent regulatory requirements for data retention and disaster recovery. Cloud disaster recovery solutions can help organizations meet these compliance requirements by providing reliable backup and recovery capabilities.

Understanding Cloudera Public Cloud implementation

You can run most of the resources in Cloudera Public Cloud such as environment, Data Lake, Data Hub, and data services and their contained services in a single replica, high availability (HA), or multiple availability zones (Multi-AZ) HA modes depending on the use case and their availability requirements.

Note: Before you deploy your use case, consider the following points:

- HA deployments require more nodes.
- Multi-AZ deployments might incur inter-region networking costs from the cloud service provider.
- There might be potential performance latency issues.

Availability zones and regions for Cloudera Public Cloud

Cloudera Public Cloud consists of two components, that is Cloudera Control Plane and Cloudera workload clusters.

Cloudera Control Plane regions

Cloudera fully operates and manages the Cloudera Public Cloud Control Plane, and provides it as a service to all of the customers in a multi-tenant manner. Cloudera Control Plane is deployed in a few separate geographic regions to satisfy the latency and data residency requirements.

As of late 2023, the following Control Plane regions are available:

Regions	us-west-1	eu-1	ap-1
Location	United States	Germany	Australia
Year opened	2019	2021	2021

Every Cloudera Account, also called a tenant, belongs to a single Cloudera Control Plane region and the account data and metadata are kept within the specified geographic boundaries. The Cloudera Control Plane regions are isolated from each other and the Cloudera accounts cannot access data from the other Control Plane regions.

Cloudera Control Plane runs in a Public Cloud context, that is it provides redundant hardware, multiple availability zone support, fast failover capabilities, and high availability Service Level Agreement (SLA). Cloudera is responsible for all the high availability, data durability, and disaster-recovery policies for the Control Plane.

Cloudera Public Cloud workload, regions, and availability zones

Availability Zone (AZ)

Contains one or more data centers. These data centers are provided with redundant power, cooling, and networking design. Typically, the availability zones are located close enough to each other to provide a very low latency (<2ms) between them in the same region. An availability zone is a failure domain, that is it might get isolated or unavailable from other AZs in the region due to malicious software deployments or natural disaster events such as earthquakes or tornadoes.

By default, Cloudera Public Cloud provisions the Data Lake, FreeIPA, and the Data Hubs in a single AWS or Azure availability zone, but you can optionally choose to deploy them across multiple availability zones (multi-AZ). It is possible to enable it for all the components or for just a few components.

Region

Collection of a minimum of three availability zones. A region is also a physical location and an independent geographic area. All availability zones within a region share low latency, high-bandwidth dedicated networking. The availability zones within a region are physically isolated and separated by meaningful distances. However, natural disasters or other events might still cause region-wide outages.

Cloudera supports more than 80 regions in three different cloud providers which include AWS region,Azure region, and GCP region.

Cloudera workloads

Run on VPCs / VNets in your cloud accounts in AWS, Azure, or GCP. When you create a Cloudera environment in the Cloudera Management Console, you must specify the region, and then create your Cloudera workloads in the Cloudera environment. You can create workloads in any of the regional Cloudera Control Planes (US, APAC, or EU).

Note: The Cloudera workload resources which include environment, Data Lake, Data Hub, or the data services are always contained within a single Public Cloud region and are not available as multi-region services.

By default, Cloudera Public Cloud provisions the environment, Data Lake, and the Data Hubs in a single AWS or Azure availability zone, but you can optionally choose to deploy them across multiple availability zones (multi-AZ). However not all Data Hub templates can be deployed across multiple availability zones.

For more information about availability zones and regions, see AWS's global infrastructure regions & AZs, Azure's availability zones and GCP's geography and regions.

AWS and Azure Multi-AZ for environments and Data Lakes

By default, Cloudera provisions Data Lake, FreeIPA, and Data Hubs in a single AWS or Azure availability zone (AZ). You can choose to deploy them across multiple availability zones (multi-AZ). It is possible to enable it for all the components or a few components.

You can choose to deploy your Cloudera environment (FreeIPA and Data Lake) and Data Hubs across multiple subnets and availability zones, where you can deploy each of these components across three or more availability zones to provide high availability and fault tolerance. The following diagram illustrates this scenario where each of the components (FreeIPA, Data Lake, and Data Hubs) are deployed across multiple subnets and multiple availability zones (*eu-central-1a*, *eu-central-1b*, and *eu-central-1c*):

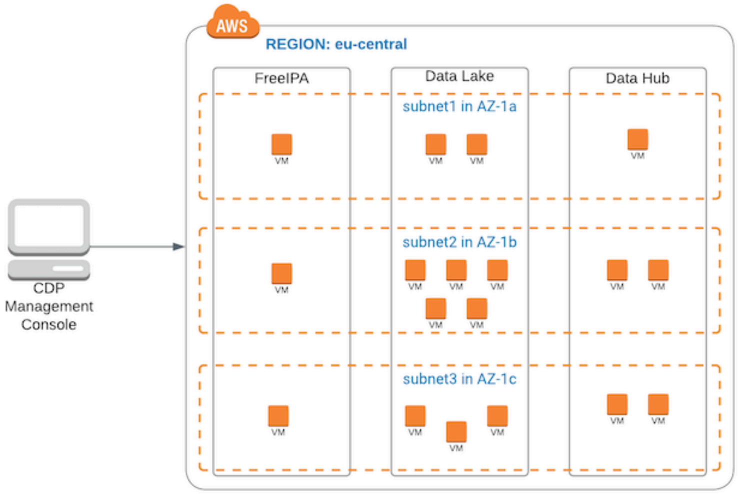

For more information, see Deploying Cloudera in multiple AWS availability zones

AWS and Azure Multi-AZ for Data Hubs

By default, Cloudera provisions Data Hubs in a single AWS or Azure availability zone (AZ), but you can optionally choose to deploy them across multiple availability zones (multi-AZ).

Data Hub clusters can be created from multiple different templates like “Data Engineering” and “Data Mart”. You can also create custom templates to deploy a custom set of supported components such as Yarn, Spark, Kafka, HBase, and so on. If your use case requires it to be zone-redundant and handle AZ failures, you can use an HA cluster template and then use the Multi-AZ flag when you create the Data Hub clusters. For more information, refer to the AWS and Azure use case documents for limitations.

In a multiple availability zones (multi-AZ) setup, the Control Plane and the Data Hub service distribute each individual host across a host group across multiple AZs. For some services like Zookeeper, the service instances are distributed across multiple AZs regardless of the host group definitions.

High availability in Cloudera Public Cloud

Cloudera workload clusters can be deployed in high availability (HA) mode to handle most of the node-level failures automatically and transparently, and to ensure that the services remain available.

The following sections explain the HA setup for Cloudera environments which contain FreeIPA instances; Cloudera Data Lakes which contain security and governance services like Ranger, RAZ, Atlas and HMS; and the Data Hubs which contain compute services like Spark, Kafka, HBase, YARN, and so on.

Tip:

- COD provides HA support. For information about COD HA, see *High Availability (Multi-AZ) for Cloudera Operational Database* in Cloudera Blog.
- Data is stored on the cloud providers’ infrastructure, therefore you can use the cloud vendor's storage HA feature.

Environment HA

FreeIPA is an open-source product that combines four identity management capabilities, that is LDAP directory, Kerberos KDC, DNS server, and Certificate Authority (CA). A FreeIPA cluster can be provisioned with a maximum of three nodes to provide high availability (HA) for Cloudera environments.

Tip: When you register a Cloudera environment using CDP CLI, you can select the node count manually. You can choose to create one to three nodes. If you do not specify the node count, Cloudera automatically configures FreeIPA with only one node.

The node provisioning depends on the Data Lake you choose for your Cloudera Public Cloud deployment.

- Enterprise Data Lake has a FreeIPA cluster with three nodes (HA) provisioned.

  Important: Cloudera recommends that the production setup always have Enterprise Data Lake configuration and run with three FreeIPA nodes.

- Light Duty Data Lake has a FreeIPA cluster with 2 nodes provisioned.

Note: Medium Data Lake for AWS, Azure, and GCP is deprecated as of Cloudera Runtime 7.2.17. Cloudera recommends all the production workloads use *Enterprise Data Lake* which is a super set of the deprecated Medium Duty Data Lake option.Medium Duty Data Lake has a FreeIPA cluster with three nodes (HA) provisioned.

FreeIPA HA cannot be added to an existing environment, only configured during environment creation. When you configure FreeIPA HA, the HA process allows an automatic failover when a FreeIPA instance fails, and then initiates a scripted manual process to recover the system with no downtime.

To understand FreeIPA, see FreeIPA identity management. For information about FreeIPA HA, see Managing FreeIPA.

Data Lake HA

Enterprise Data Lake offers HA and fault resilience for most components inside the Data Lake whereas Light duty Data Lake does not provide HA.

Cloudera Public Cloud offers the following types of Data Lakes depending on the scale and production-level requirements for the workload:

Light Duty Data Lake

This Data Lake is not highly available and is not suitable for production workloads. However, you can use it for demonstration and proof of concept (POC) purposes, and during early evaluation phases.

Enterprise Data Lakes

Enterprise Data Lake offers HA and fault resilience and the configuration enables zero downtime upgrade (ZDU) for most of the components inside the Data Lake. Enterprise Data Lakes are available as of 7.2.17 for new environments only.

Availability of services depends on the node being repaired. With the exception of the gateway and auxiliary nodes, the remaining groups can typically survive a single node failure without affecting the workloads or UI/API access.

In the event of a gateway node failure on an Enterprise Data Lake, the load-balancer seamlessly routes to the other gateway node.

For more information, see Data Lake scale.

Data Hub HA

The Cloudera Data Hub service allows you to create workload clusters to run different components like Spark, Kafka, HBase, Impala, Hive, Nifi and so on.

You can create a cluster from a predefined or a custom cluster template. A cluster template is a declarative definition of a cluster that defines the cluster topology which includes the cluster host groups, and all the cluster services and their components running on them. Data Hub provides default cluster templates, and it also allows you to upload your own cluster templates (also called custom cluster template).

In production setups, Cloudera recommends that you use templates that are marked *High Availability*. You can identify these templates by the term *HA* in the template name.

Some Data Hubs and all the HA templates use external databases. Multi-AZ templates and custom templates create databases in a HA multi-AZ setup.

The following table lists the Data Hub templates with its corresponding components and services that support HA:

Data Hub template	Component	Services
Enterprise SDX template with Atlas, HMS, Ranger and other services they are dependent on. Services like HDFS, HBASE, RANGER, HMS have HA.	Hive	hive-HIVEMETASTORE
	Ranger	• ranger-RANGER_ADMIN • RAZ
	HDFS	• hdfs-NAMENODE • hdfs-FAILOVERCONTROLLER • hdfs-DATANODE • hdfs-JOURNALNODE
	Atlas	atlas-ATLAS_SERVER
	Zookeeper	zookeeper-SERVER
	Knox	• knox-KNOX_GATEWAY • knox-IDBROKER
	Kafka	• kafka-KAFKA_BROKER
	Solr	• solr-SOLR_SERVER
	HBase	• hbase-REGIONSERVER • hbase-MASTER
Data Engineering HA	HDFS	• hdfs-DATANODE • hdfs-FAILOVERCONTROLLER • hdfs-HTTPFS • hdfs-JOURNALNODE • hdfs-NAMENODE
	Zookeeper	zookeeper-SERVER
	Hive	• hive-HIVESERVER2 • hms-HIVEMETASTORE • hue-HUE_LOAD_BALANCER • hue-HUE_SERVER
	Spark	• livy-LIVY_SERVER • spark_on_yarn-SPARK_YARN_HISTORY_SERVER
	YARN	• yarn-NODEMANAGER • yarn-NODEMANAGER-COMPUTE • yarn-RESOURCEMANAGER
	Oozie	oozie-OOZIE_SERVER
	Knox	knox-KNOX
Data Engineering - Spark3 HA	HDFS	• hdfs-DATANODE • hdfs-FAILOVERCONTROLLER • hdfs-HTTPFS • hdfs-JOURNALNODE • hdfs-NAMENODE
	Zookeeper	zookeeper-SERVER
	Hive	• hive-HIVESERVER2 • hms-HIVEMETASTORE • hue-HUE_LOAD_BALANCER • hue-HUE_SERVER

Data Hub template	Component	Services
	Spark	• livy-LIVY_SERVER • spark_on_yarn-SPARK_YARN_HISTORY_SERVER
	YARN	• yarn-NODEMANAGER • yarn-NODEMANAGER-COMPUTE • yarn-RESOURCEMANAGER
	Oozie	oozie-OOZIE_SERVER
	Knox	knox-KNOX
	Spark	• livy_for_spark3-LIVY_SERVER • spark3_on_yarn-SPARK_YARN_HISTORY_SERVER
Streams Messaging High Availability with Apache Kafka, Schema Registry, Streams Messaging Manager, Streams Replication Manager, and Cruise Control	Streaming	• schemaregistry-SCHEMA_REGISTRY_SERVER (used for topic schema management for governance) • kafka-KAFKA_BROKER • streams_replication_manager-STREAMS_REPLICATION_MANAGER_SERVICE • streams_replication_manager-STREAMS_REPLICATION_MANAGER_DRIVER • kafka-KAFKA_CONNECT (connector for external database/NiFi to Kafka) • kafka-KAFKA_KRAFT (in technical preview) replaces zk for kafka operation)
	Zookeeper	zookeeper-SERVER
	Knox	knox-KNOX (used for schema registry and Replication Manager)
Real-time Data Mart - Apache Impala, Hue, Apache Kudu,and Apache Spark	-	• kudu-MASTER • yarn-NODEMANAGER • impala-IMPALAD-EXECUTOR • kudu-TSERVER
Flow Management Heavy Duty with Apache NiFi, Apache NiFi Registry, and Schema Registry	NiFi	nifi-NIFI_NODE
Streaming Analytics Heavy Duty with Apache Flink	-	• zookeeper-SERVER • hdfs-FAILOVERCONTROLLER • hdfs-JOURNALNODE • hdfs-NAMENODE • kafka-KAFKA_BROKER • yarn-RESOURCEMANAGER • hdfs-DATANODE • yarn-NODEMANAGER
Data Mart with Apache Impala and Hue	Impala	• impala-IMPALAD-EXECUTOR