

High availability in Cloudera Data Engineering

Date published: 2020-07-30

Date modified: 2025-07-30

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps. The letter 'E' is stylized with a horizontal bar that extends to the right and then turns downwards, forming a unique graphic element.

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

Multiple Availability Zone support and limitations in Cloudera Data Engineering.....	4
---	----------

Multiple Availability Zone support and limitations in Cloudera Data Engineering

This document highlights the limitations of Cloudera Data Engineering's support for multiple Availability Zones.

Availability Zones overview

Each AWS or Azure Region is subdivided into separate Availability Zones (AZs), which have their own power, cooling, and network connectivity, forming an isolated failure domain. If you run workloads in multiple AZs, you can ensure that your applications withstand even a complete AZ failure.

Services and tools

This section lists the configuration, the behavior during failures, and potential limitations of:

- Elastic Kubernetes Service (EKS)
- Relational Database Service (RDS)
- AWS Elastic Compute Cloud (EC2)
- Load balancers
- Elastic File System (EFS)

Elastic Kubernetes Service (EKS)

- The EKS control plane is designed for High Availability (HA). It consists of at least two API server instances and three ETCD instances distributed across three AZs within an AWS Region.
- EKS automatically detects and replaces unhealthy control plane instances, restarting them as needed across different AZs to maintain availability and reliability.

For more information, see [Understand resilience in Amazon EKS clusters](#) in the AWS documentation.

Relational Database Service (RDS)

- Cloudera Data Engineering supports RDS creation in multiple AZs.
- RDS provides a Reboot with failover option: it starts a standby instance in case of a failure. This causes a downtime of approximately 2 to 3 minutes.
- Multiple AZ RDS creation is not available through the CLI or the UI; you can only perform it through the Cloudera Data Engineering API.

AWS Elastic Compute Cloud (EC2)

- EC2 instances for Cloudera Data Engineering can leverage multiple subnets across AZs during service creation, creating Auto Scaling Groups (ASGs) with multiple AZ capability.
- In case of an AZ failure, instances start in other AZs if configured so.

Load balancers

Load balancers, similar to EC2 instances, support multiple subnets, making them resilient to AZ failures.

Elastic File System (EFS)

- Cloudera Data Engineering provides EFS with regional availability, enabling mount targets across multiple AZs.
- EFS is highly available and resilient to AZ failures.

Cloudera Data Engineering components

Most Cloudera Data Engineering components, such as the Airflow server, the Jobs API server, and others, are deployed using Kubernetes deployments. In the event of a node failure, these components are capable of self-healing: starting on newly provisioned nodes.

The downtime for these services depends on the cumulative time required for the nodes, pods, and the RDS instance to start and become operational.

Cloudera Data Engineering jobs and services behavior during idbroker node failures

Cloudera Data Engineering jobs remain resilient during an AZ failure. After the failure, the Spark driver reroutes to an alternate region URL.

Cloudera Data Engineering Spark jobs behavior during Spark driver AZ failure

From an RDS perspective, Cloudera Data Engineering jobs, Spark, and Airflow remain resilient, successfully enduring multiple reboots in both single and multiple AZ modes.

Limitations

Cloudera has not performed an end-to-end testing of the AZ failure scenario. The initial analysis indicates that Cloudera Data Engineering cannot fully support multiple AZ.

Spark driver is not resilient to AZ failure

Apache Spark jobs might fail if the Spark driver pod is stopped during a node failure, as the Spark driver maintains the job state and coordination. In such cases, you need to restart the job manually or leverage retry mechanisms using Apache Airflow. However, when employing the retry mechanism, you need to wait until the API server and Livy restarts. Cloudera needs to further investigate and test multiple AZ failure scenarios using tools, such as [Amazon FIS](#). For more information, see [Simulating Kubernetes-workload AZ failures with AWS Fault Injection Simulator](#) in the AWS documentation.

Cloudera Data Engineering jobs and services, especially tgt generation, behavior during FreeIPA node failures

- Single AZ failure: jobs run seamlessly in the event of a single AZ failure.
- Dual AZ failure: when two AZs fail simultaneously, the following issues arise:
 - Job submissions temporarily fail due to DNS caching.
 - While new job submissions eventually succeed, jobs remain stuck in the STARTING phase.
 - Tgtgen pods fail with connection errors.

Recovery: recovery occurs once the affected AZs are restored. However, jobs that were stalled during the failure do not transition back to RUNNING state.

Limitation of AWS Elastic Compute Cloud (EC2)

The Cloudera shared infrastructure node group is currently restricted to a single AZ, due to historical reasons. The YuniKorn scheduler is created without multiple AZ support, and if the AZ with YuniKorn fails, the cluster may become inoperable.

While it is possible to override this limitation within the Cloudera Data Engineering code during cluster requests by creating the Cloudera shared infrastructure node group ASG directly, Cloudera does not recommend it due to the complexities involved with specifying defaults.

Data Lake failover

As Cloudera Data Lake supports multiple AZ, if the endpoints of any of the services change during failover, Cloudera Data Engineering might not be able to adapt to the new endpoints. For example, tgtGen may not be able to fetch tgt if the Data Lake fails over. The tgt service is responsible for generating Kerberos tickets. One particular scenario is if the CM node (Data Lake node managed by Cloudera Manager) falls into the failed AZ, the CM is unhealthy, which results in failed tgt generation.

Conclusion

While the individual cloud components are compatible with multiple AZs and can fail over, there are limitations, such as Spark pods, Cloudera shared infrastructure node group, and Data Lake connections that do not work as expected in failover cases.