

# Deploying CDP in Multiple AWS Availability Zones (Preview)

Date published: 2021-12-07

Date modified: 2022-04-01

## Legal Notice

© Cloudera Inc. 2021. 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.

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.

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

## Contents

<b>Legal Notice</b>	<b>2</b>
<b>Contents</b>	<b>3</b>
<b>Introduction</b>	<b>4</b>
Single-AZ vs multi-AZ deployment	4
Use cases	6
Limitations	6
<b>Enabling multi-AZ</b>	<b>7</b>
Install Beta CDP CLI	7
Register a multi-AZ environment	8
CDP web interface	8
CDP CLI	8
Create a multi-AZ Data Hub	9
CDP web interface	10
CDP CLI	10
<b>Scaling a multi-AZ Data Hub</b>	<b>11</b>

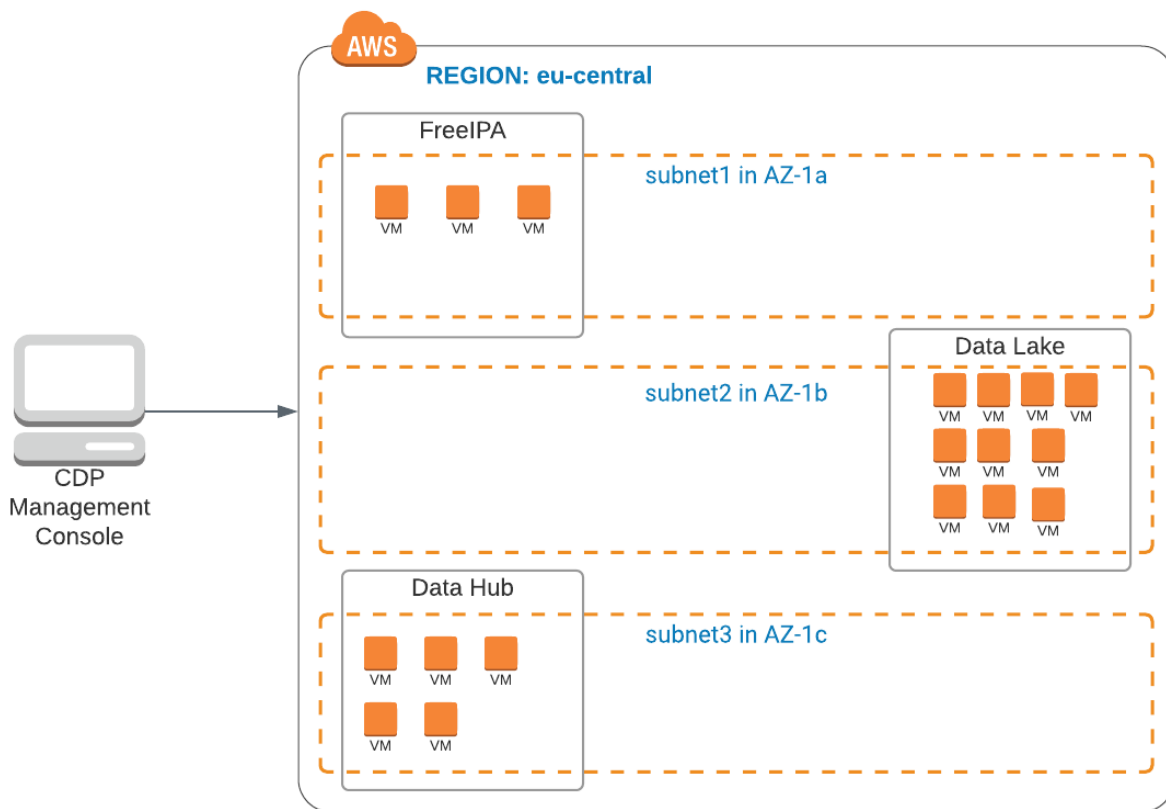
*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

# Introduction

By default, CDP provisions Data Lake, FreeIPA and Data Hubs in a single Availability Zone (AZ), but you can optionally choose to deploy them across multiple Availability Zones (multi-AZ). It is possible to enable it either for all or some of these components.

## Single-AZ vs multi-AZ deployment

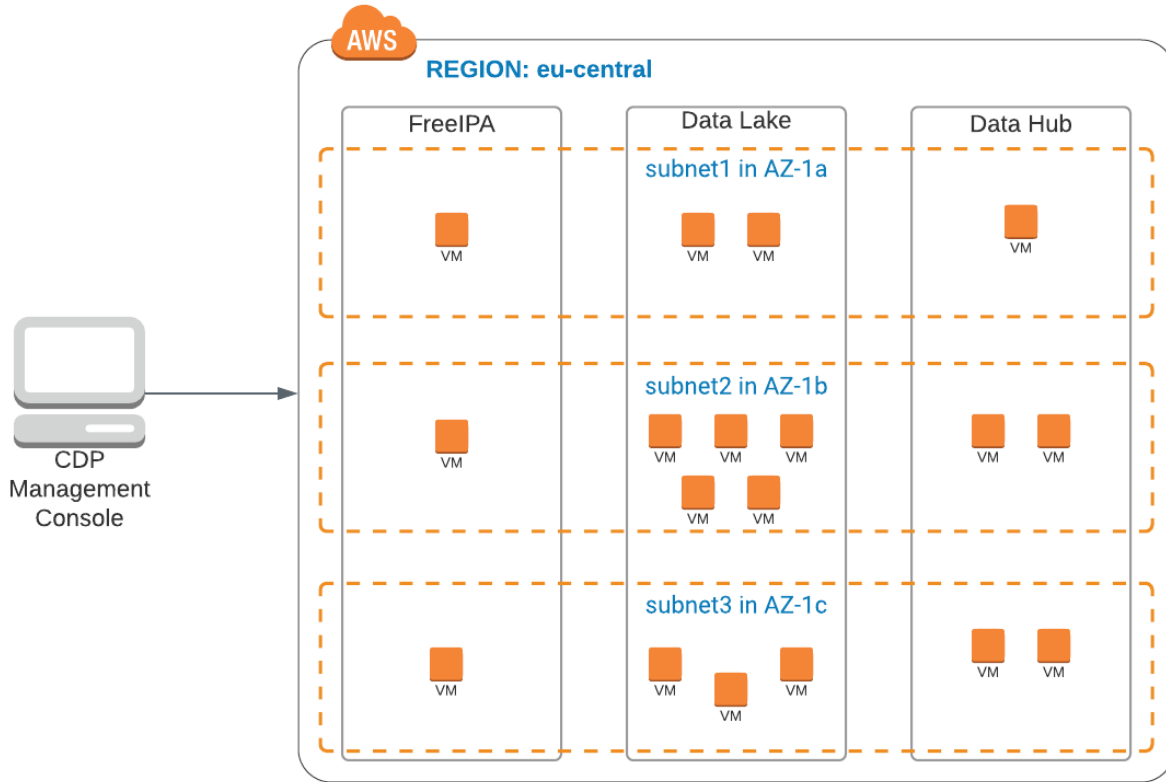
Each AWS region has multiple availability zones, which act as failure domains, preventing small outages from affecting entire regions. By default, a CDP environment (FreeIPA and Data Lake) and Data Hubs running in it are deployed across one subnet, meaning that each of them has VMs spread across a single availability zone (as on AWS every subnet is related to a single availability zone). This is illustrated in the following diagram, where all of the resources run in the eu-central region, and each of them is in a different subnet and availability zone (eu-central-1a for the FreeIPA, eu-central-1b for the Data Lake, and eu-central-1c for the Data Hubs):



If you choose to deploy your CDP environment (FreeIPA and Data Lake) and Data Hubs across multiple subnets and availability zones, each of these components is spread across three or

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

more availability zones, providing high availability and fault tolerance. This is illustrated in the following diagram, where each of the components (FreeIPA, Data Lake, and Data Hubs) is deployed across multiple subnets and multiple availability zones: eu-central-1a, eu-central-1b, and eu-central-1c:



The Data Lake deployed in this setup is the Medium Duty Data Lake, which has enough nodes to utilize multiple availability zones for providing high availability and fault tolerance.

Multi-AZ can be enabled for Data Lake, FreeIPA, and Data Hubs. It is possible to enable it either for all or some of these components.

Instances within a host group are distributed across availability zones equally, with the least used availability zone(s) preferred. For example, if there is a group of instances across three availability zones with 100 instances in AZ-1, 30 in AZ-2 and 30 in AZ-3, when an upscale of 10 nodes is requested, CDP provisions 5 instances in AZ-2 and 5 in AZ-3.

When a failure happens and a cluster needs to be repaired, the replacement VMs are always provisioned in the same subnet and availability zone as the old ones since the detached disks can only be reattached to a VM in the same availability zone. This means that if there is an availability zone outage, cluster repair is not possible.

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

## Use cases

A multi-AZ Data Lake and FreeIPA constitute a resilient environment that provides a solid basis for multi-AZ Data Hubs and CDP data services. Data Hubs and CDP data services depend on the FreeIPA instance in the Data Lake to provide DNS resolution. Deploying FreeIPA across multiple availability zones ensures that critical DNS resolution is available in the event of an availability zone outage. Furthermore, a Medium Duty Data Lake provides high availability, and additional compute and memory resources for key SDX services and is recommended for production workloads.

Deploying your Data Hubs across multiple availability zones is key if your mission-critical applications depend on HBase and Kafka. Multiple availability zone deployment for operational workloads is considered best practice by the cloud vendors. It ensures that your applications can continue to run in the event of a single availability zone outage.

When an entire availability zone fails, HBase automatically rebalances regions among the remaining instances in the cluster to maintain availability. The write-ahead log (WAL), which is replicated across the three availability zones is automatically replayed by the newly assigned region servers in other availability zones to ensure writes to the database are not lost.

When using the multi availability zone feature, CDP ensures that Kafka replicates partitions across brokers in different availability zones. During an availability zone failure this ensures that no data is lost and applications can continue to access the data they need. Cruise Control, which is deployed alongside every Kafka cluster in CDP Public Cloud, detects that topics need to be rebalanced to the remaining brokers. Once the availability zone is back online, you can repair your Kafka cluster, restoring the initial broker distribution across availability zones. Afterwards Cruise Control kicks in and ensures that all topic partitions are balanced across the cluster.

## Limitations

The following limitations apply when deploying a multi-AZ CDP:

- When an AZ is down you cannot create a new Data Hub, and create or activate CDP data services within the environment. Existing workloads will continue to work.
- When an AZ is down you cannot stop and restart Data Hubs.
- When an AZ is down, existing running workloads will continue to run. Some Data Lake services will still be functioning (because they run across more than one AZ) but others may be down if the AZ where they are running is down. The following table shows which services aren't affected by an AZ outage (column 1 and 2) and those that may potentially be affected (column 3):

<b>Services running in all available AZs</b>	<b>Services running in two AZs</b>	<b>Services running in a single AZ</b>
Datanode	Atlas Server	Cloudera Manager

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

HBase Gateway HBase Master HBase RegionServer HDFS Gateway HDFS JournalNode Hive Gateway Kafka Broker Kafka Gateway Solr Server	HDFS Failover Controller Hive Metastore IDBroker Knox Gateway Namenode Ranger Admin Solr Gateway Zookeeper Server	Ranger Usersync Ranger Tagsync
---------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------

- While CDP CLI allows you to specify multiple subnets per Data Hub’s host group, due to load balancer requirements only one subnet per AZ can be configured for the Cloudera Manager node group of the Data Hub cluster.
- While by default CDP deploys Data Lake, FreeIPA, and Data Hubs using CloudFormation, in case of the multi-AZ setup no CloudFormation is used. This affects vertical scaling: After vertical scaling (for example during cluster repair or upgrade) CDP always provisions the original instance types. For example, if you are using vertical scaling runbooks, repair and upgrade operations will use the previous instance type and root volume settings.

## Enabling multi-AZ

Multi-AZ can be enabled for Data Lake, FreeIPA, and Data Hubs. It is possible to enable it either for all or some of these components.

To enable multi-AZ, you should:

- If you would like to use CLI to provision multi-AZ environments and/or Data Hubs, install Beta CDP CLI. The Beta CDP CLI includes a few customization options for Data Hub creation and scaling that are not available via CDP web UI.
- Register a multi-AZ CDP environment with a Medium Duty Data Lake.
- Create a multi-AZ Data Hub within an existing environment.

Detailed steps are provided below.

## Install Beta CDP CLI

In order to deploy CDP into multiple availability zones on AWS via CLI, you need to download and install Beta CDP CLI.

### Prerequisites

Do not install both the standard and beta CLIs in the same Python environment, as they use the same entry points and will therefore conflict. Either use a separate virtual environment or uninstall the cdpcli first before installing cdpcli-beta.

### Steps

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided ‘as is’ without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

1. Install Python if needed, as described in [CDP CLI documentation](#) for your platform. Do not install CDP CLI. Instead, proceed to step 2.
2. Run the following to install Beta CDP CLI:  

```
pip3 install cdpcli-beta
```
3. Configure access keys as described in [Generating an API access key](#).
4. If you have previously installed Beta CDP CLI, you can update it to the latest version using:  

```
pip3 install cdpcli-beta --upgrade --user
```


## Register a multi-AZ environment

You can register a multi-AZ environment via CDP UI or Beta CDP CLI.

### CDP web interface

Register your environment as usual, just make sure to do the following:

1. On the **Data Access and Data Lake Scaling** page, select the Medium Duty Data Lake.
2. On the same page, enable the following option in order to enable multi-AZ for Data Lake:



The option is disabled by default. The option only appears when a Medium Duty Data Lake is selected.

3. On the **Region, Networking, and Security** page, enable the following option in order to enable multi-AZ for FreeIPA:



The option is disabled by default.

**Note:** You may choose to enable multi-AZ for Data Lake only or for FreeIPA only. There is no requirement to enable both.

### CDP CLI

Register your AWS environment via CDP CLI as usual, but when creating a template for the environment registration commands, add the following to enable multi-AZ:

- In the environment creation command, add the `multiAz=true` to the `--free-ipa` option.
- In the Data Lake creation command, add the `--multi-az` option.

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*



- In the Data Lake creation command, make sure to specify the Medium Duty Data Lake.

The following example illustrates these necessary additions, which are emphasized using bold typeface:

```

cdp environments create-aws-environment \
--environment-name tb-multiaz-env \
--credential-name tb-cldr-acc \
--region "eu-central-1" \
--security-access cidr=0.0.0.0/0 \
--no-enable-tunnel \
--authentication publicKeyId="test" \
--log-storage
storageLocationBase=s3a://cb-group/test,instanceProfile=arn:aws:iam::152413716728:instance-profile/mock-idbroker-admin-role \
--vpc-id vpc-03e505817e3619238 \
--subnet-ids subnet-013855b2gc32c2cd8 subnet-02b9054ef829374fe
subnet-085c9ff36b38c0b35 \
--free-ipa instanceCountByGroup=3,multiAz=true

```

```

cdp environments set-id-broker-mappings \
--environment-name tb-multiaz-env \
--data-access-role arn:aws:iam::152413716728:role/mock-idbroker-admin-role \
--ranger-audit-role arn:aws:iam::152413716728:role/mock-idbroker-admin-role \
--set-empty-mappings

```

```

cdp datalake create-aws-datalake \
--datalake-name tb-multiaz-dl \
--environment-name tb-multiaz-env \
--cloud-provider-configuration
instanceProfile=arn:aws:iam::152413716728:instance-profile/mock-idbroker-assumer
,storageBucketLocation=s3a://cb-group/test \
--scale MEDIUM_DUTY_HA \
--runtime 7.2.12 \
--no-enable-ranger-raz \
--profile mowdev \
--multi-az

```

## Create a multi-AZ Data Hub

You can create a multi-AZ Data Hub via CDP UI or Beta CDP CLI within an environment. Note that the CLI allows you to manually specify subnets, which is not possible via the UI.

## CDP web interface

Create your Data Hub as usual. In the **Advanced Options > Network and Availability**, you select the multiple subnets across which the Data Hub is to be provisioned. If multiple subnets are selected, the Cloudera Manager node group will only have one subnet for each AZ; All other nodes will have all the selected subnets.

## CDP CLI

When creating a multi-AZ Data Hub, Data Hub creation request should contain the `--multi-az` option. For example:

```

cdp datahub create-aws-cluster \
--cluster-name tb-datamart-multiaz \
--environment-name tb-multiaz-env \
--cluster-template-name "7.2.12 - Data Mart: Apache Impala, Hue" \
--instance-groups
nodeCount=3,instanceGroupName=coordinator,instanceGroupType=CORE,instanceType=r5
d.4xlarge,rootVolumeSize=100,attachedVolumeConfiguration=[\{\volumeSize=300,volu
meCount=2,volumeType=ephemeral\}\],recoveryMode=MANUAL,volumeEncryption=\{enable
Encryption=false\}
nodeCount=1,instanceGroupName=master,instanceGroupType=GATEWAY,instanceType=r5.4
xlarge,rootVolumeSize=100,attachedVolumeConfiguration=[\{\volumeSize=100,volumeC
ount=1,volumeType=standard\}\],recoveryMode=MANUAL,volumeEncryption=\{enableEncr
yption=false\}
nodeCount=2,instanceGroupName=executor,instanceGroupType=CORE,instanceType=r5d.4
xlarge,rootVolumeSize=100,attachedVolumeConfiguration=[\{\volumeSize=300,volumeC
ount=2,volumeType=ephemeral\}\],recoveryMode=MANUAL,volumeEncryption=\{enableEnc
ryption=false\} \
--image id=23e20852-7865-4980-a045-539296340b55,catalogName=cdp-default \
--profile mowdev \
--multi-az

```

By default, Data Hub instances are distributed across the subnets provided during environment registration. If you prefer to specify the subnets manually, you can provide `--subnet-ids` option and specify subnets where you would like to deploy the hosts, overwriting the default node distribution. In this case, add the following to the Data Hub creation command:

```

--subnet-ids "subnet-013855b2fc32c2cd8" "subnet-02b9054ec829374fe"
"subnet-085c9ff36b38c0b35"

```

If multiple subnets are provided, the Cloudera Manager node group will only have one subnet for each AZ; All other nodes will have all the selected subnets.

*This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.*

## Scaling a multi-AZ Data Hub

If there is an availability zone that is offline, CDP may not detect the outage. In such a case, if you know that a certain availability zone is offline, you can scale your cluster and manually specify where the new nodes should be provisioned.

When scaling a multi-AZ cluster, CDP automatically distributes the new nodes in a round-robin fashion across all available availability zones, prioritizing the least used availability zones. If you prefer to manually control the distribution of nodes across subnets during Data Hub scaling, the desired availability zones can be controlled via the related subnets during upscales with the optional `--preferred-subnet-ids` field.

For example:

```
cdp datahub scale-cluster --profile mowdev \
--cluster-name tb-datamart-multiaz\
--instance-group-name "coordinator"
--instance-group-desired-count 5
--preferred-subnet-ids "subnet-013855b2fc32c2cd8" "subnet-02b9054ec829374fe"
"subnet-085c9ff36b38c0b35"
```

If you manually specify the subnets in this manner, this overwrites the default behavior.