

Working with data lakes 2

Working with Data Lakes (TP)

Date of Publish: 2019-02-06



<https://docs.hortonworks.com/>

Contents

Data lake overview.....	3
Data lake blueprints.....	5
Setting up a data lake.....	5
Prerequisites.....	6
Register databases and LDAP.....	6
Create a data lake.....	7
Create attached workloads.....	10
Performing manual HA data lake recovery.....	11

Data lake overview

Cloudbreak allows you to deploy a long-running data lake instance and attach it to multiple ephemeral workload clusters.



Note:

This feature is technical preview: It is not suitable for production.



Note:

The data lake deployment option is only suitable for AWS, Azure, and Google; It is not suitable for OpenStack.

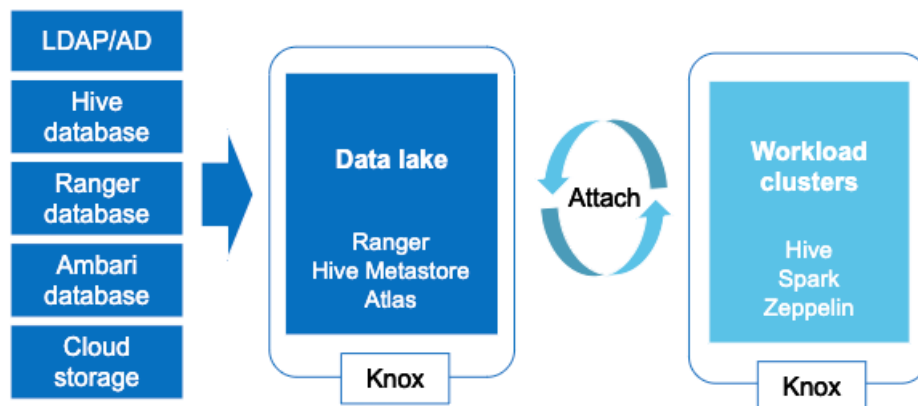
What is a data lake

A data lake provides a way for you to centrally apply and enforce authentication, authorization, and audit policies across multiple ephemeral workload clusters.

“Attaching” your workload cluster to the data lake instance allows the attached cluster workloads to access data and run in the security context provided by the data lake.

While workloads are temporary, the security policies around your data schema are long-running and shared for all workloads. As your workloads come and go, the data lake instance lives on, providing consistent and available security policy definitions that are available for current and future ephemeral workloads. All information related to schema (Hive), security policies (Ranger), and audit (Ranger) is stored on external locations (external databases and cloud storage). Furthermore, Ambari database is external and can be recovered.

This is illustrated in the following diagram:



Once you’ve created a data lake instance, you have an option to attach it to one or more ephemeral clusters. This allows you to apply the authentication, authorization, and audit across multiple workload clusters.

Basic terminology

The following table explains basic data lake terminology:

Term	Description
Data lake	Runs Ranger, which is used for configuring authorization policies and is used for audit capture. Runs Hive Metastore, which is used for data schema.
Workload clusters	The clusters that get attached to the data lake to run workloads. This is where you run workloads such as Hive via JDBC.

Components

The following table explains the components of a data lake:

Component	Technology	Description
Schema	Apache Hive	Provides Hive schema (tables, views, and so on). If you have two or more workloads accessing the same Hive data, you need to share schema across these workloads.
Policy	Apache Ranger	Defines security policies around Hive schema. If you have two or more users accessing the same data, you need security policies to be consistently available and enforced.
Audit	Apache Ranger	Audits user access and captures data access activity for the workloads.
Governance	Apache Atlas	Provides metadata management and governance capabilities.
Directory	LDAP/AD	Provides an authentication source for users and a definition of groups for authorization.
Gateway	Apache Knox	Supports a single workload endpoint that can be protected with SSL and enabled for authentication to access to resources.




Note:

Apache Atlas is not available as part of the HA blueprint but is available as part of a separate blueprint.

Data lake high availability

If you select to use the HA blueprint, your data lake services will run in HA mode. Data lake HA blueprint deploys a three-node cluster by default:



The exact number of nodes can be modified, but odd node number is HA best practice. The exact number of nodes can be modified, but odd node number is HA best practice. Since in a high availability data lake Ambari database must in all cases be external, it is available even if the Ambari server host goes down.

Cloudbreak monitors high availability data lakes, ensuring that when host-level failures occur, they are quickly resolved by deleting and replacing failed nodes of the high availability data lake (“HA data lake”). Once a failure is reported, it is repaired automatically (if enable auto recovery is selected), or options are available for you to repair the failure manually (if enable auto recovery is not selected).

When a data lake is launched, one of the hosts is selected as the primary gateway host, whereas the remaining hosts are considered secondary gateway (also known as “core”) hosts. The primary gateway host serves as the Ambari server node (with Ambari agents are connecting to that host only) and the Knox gateway node.

Two recovery scenarios are possible:

- If one of the secondary gateway hosts goes down, a downscale and upscale flow happens to replace the missing node.
- If the primary gateway host goes down, Cloudbreak first chooses a different node to be the primary gateway and launches Ambari server on that node. The new primary gateway node connects to the external Ambari database and Ambari agents automatically rejoin the new Ambari server. Next, the downscale and upscale flow happens to replace the missing node, and the newly added node becomes a secondary gateway node.

Since Cloudbreak assigns aliases to all nodes (Based on the "Custom domain" and "Custom hostname" parameters), the new nodes are assigned the same aliases as the failed nodes.

Once the repair is complete, all workload clusters attached to the HA data lake are updated so that they can communicate with the newly added nodes of the HA data lake.

Related Information


[Using an external authentication source for clusters](#)

[Gateway Configuration](#)

Data lake blueprints

When creating a data lake, you can choose from one of the two available blueprints.

The following data lake blueprints are provided by default in Cloudbreak:

Blueprint	Description	Node count
HDP 3.1 Data Lake: Apache Ranger, Apache Hive Metastore	Includes Apache Ranger and allows all clusters attached to a data lake to connect to the same Hive Metastore.  Note: Hive Metastore has been removed from the HDP 3.x data lake blueprints, but setting up an external database allows all clusters attached to a data lake to connect to the same Hive Metastore.	Includes a single master host group and must include a single node.
HDP 2.6 Data Lake: Apache Ranger, Apache Atlas, Apache Hive Metastore	Includes Apache Ranger, Apache Atlas, and Apache Hive Metastore.	Includes a single master host group and must include a single node.
HDP 2.6 Data Lake: Apache Ranger, Apache Hive Metastore HA	Includes Apache Ranger and Apache Hive Metastore in HA mode. Automatic and manual recovery options are available for this type of data lake.	Includes two master host groups. We recommend either 3 or 5 nodes total for this type of cluster. By default the node count is 3.

Depending on your use case, select one of these blueprints.

Setting up a data lake

Setting up a data lake involves meeting the prerequisites, registering external resources in Cloudbreak, and creating a data lake. Once a data lake is running, you can create workload clusters attached to the data lake.

Refer to the following table to learn more about the data lakes prerequisites:

Step	Where to perform
Review available data lake blueprints and select one that you would like to use.	Documentation or Cloudbreak web UI
Meet the prerequisites: <ul style="list-style-type: none"> • Create an external database for Hive metastore. • Create an external database for Ranger. • If you are planning to use the HA blueprint, create an external database for Ambari. • Create an external authentication source for LDAP/AD. • Prepare a cloud storage location (depending on your cloud provider, this should be, on Amazon S3, Azure's ADLS or WASB, or GCS) for default Hive warehouse directory and Ranger audit logs). 	You must create these resources on your own, outside of Cloudbreak. You may use one database instance and create two databases.
Register the two databases and LDAP	In the Cloudbreak web UI > External Sources
Create a data lake	In the Cloudbreak web UI > Create cluster
Create clusters attached to the data lake	In the Cloudbreak web UI > Create cluster

Prerequisites

The following resources (databases, LDAP, and cloud storage locations) must be created outside of Cloudbreak prior to creating a data lake.

Steps

1. Set up two external database instances, one for the HIVE component, and one for the RANGER component. For supported databases, refer to [Supported databases](#).
2. If you are planning to use the HA blueprint, also set up an external database for the AMBARI component. For supported databases, refer to [Supported databases](#).
3. Create an LDAP instance and set up your users inside the LDAP.
4. Prepare a cloud storage location for default Hive warehouse directory and Ranger audit logs.

As an outcome of this step, you should have the external resources available. In the next step, you are required to provide the information related to these external resources to Cloudbreak.

Related Information

[Supported databases](#)

Register databases and LDAP

After creating the external resources listed in prerequisites, you must register these resources in the Cloudbreak web UI.

Steps

1. Register each of your two RDS instances created as part of the prerequisites in the Cloudbreak web UI, under External Sources > Database Configurations. For instructions, refer to [Register an external database](#).
 - When registering the database for Hive, select Type > Hive.
 - When registering the database for Ranger, select Type > Ranger.
 - If using the HA blueprint, you must also register an external database for Ambari. When registering the database for Ambari, select Type > Ambari.
2. Register your LDAP (created as part of the prerequisites) in the Cloudbreak web UI, under External Sources > Authentication Configurations. For instructions, refer to [Register an authentication source](#).

As an outcome of this step, you should have two external databases and one authentication source registered in the Cloudbreak web UI. Now you can create a data lake.

Related Information

[Register an external database](#)

[Register an authentication source](#)

Create a data lake

After meeting the prerequisites and registering the external resources in Cloudbreak, create a data lake by using the create cluster wizard.

The instructions below only list data-lake specific steps. For information on other cluster options refer to the documentation for creating clusters on your cloud platform.

Steps

1. In Cloudbreak web UI, navigate to Clusters, click on Create Cluster.
2. On the General Configuration page:
 - Under Cluster Name, provide a name for your data lake.
 - Under Cluster Type, choose one of the two available "Data Lake" blueprints: either "Data Lake: Apache Ranger, Apache Atlas, Apache Hive Metastore" or "Data Lake: Apache Ranger, Apache Hive Metastore HA".

The screenshot shows the 'General Configuration' page in the Cloudbreak web UI. The page is divided into 'BASIC' and 'ADVANCED' tabs, with 'BASIC' selected. A sidebar on the left lists configuration steps: 1. General Configuration (selected), 2. Hardware and Storage, 3. Network and Availability, 4. Cloud Storage, 5. Cluster Extensions, 6. External Sources, 7. Gateway Configuration, 8. Network Security Groups, and 9. Security. The main content area is titled 'General Configuration' and includes the instruction 'Specify basic parameters for your cluster.' The form contains the following fields: 'Select Credential' (dropdown menu set to 'aws-key-based (AWS)'), 'Cluster Name' (text input field containing 'test-data-lake'), 'Region' (dropdown menu set to 'EU (Ireland)'), and 'Platform Version' (dropdown menu set to 'HDP 2.6'). The 'Cluster Type' dropdown menu is highlighted with a green box and a green arrow pointing to it from the right. The selected option is 'Data Lake: Apache Ranger, Apache Hive Metastore', which is also highlighted with a green box. Below this option, there is a sub-option 'DATA LAKE Data Lake: Apache Ranger, Apache Hive Metastore'.

3. (Only if using the HA blueprint) On the Hardware and Storage page you can select the following for each host group:
 - a. Under Instance Count, you can optionally specify how many nodes should be included in each host group. By default, Cloudbreak creates the minimum viable number of nodes. We recommend placing an odd number of nodes in each host group. A total of 3 or 5 instances is recommended.
 - b. You can optionally select to Enable Auto Recovery. Enabling this option will allow Cloudbreak to automatically recover any failed nodes. Without checking this option, you will have to manually trigger recovery of the failed nodes.

BASIC **ADVANCED**

General Configuration

2 **Hardware and Storage**

3 Network and Availability

4 Cloud Storage

5 Cluster Extensions

6 External Sources

7 Gateway Configuration

8 Network Security Groups

9 Security

For each host group, specify the number of instances, instance type, and storage attached.

Master1 Node Ambari Server

Instance Type : m5.2xlarge (8vCPU & 32 GB Memory)

Instance Count : 1

Storage Type : Magnetic

Encryption : Not encrypted

Attached Volume per Instances : 1

Volume Size : 100 GB

Root Volume Size : 50

Use Spot Instances : No

Enable Auto Recovery : No

Master2 Nodes Ambari Server

Instance Type : m5.2xlarge (8vCPU & 32 GB Memory)

Instance Count : 2

4. (Only if using the HA blueprint) On the Network and Availability page, enter:

- a. Custom Domain: Enter some domain name that Cloudbreak can use locally. For example “mydatalake.local”. This domain name is for local use and does not require DNS services.
- b. Custom Hostname: Enter some name convention to use for the host names. For example “prod”.

For example, if using “mydatalake.local” as a custom domain and “prod” as a host name, the actual host names will be prod0.<cluster-name>.mydatalake.local, prod1.<cluster-name>.mydatalake.local, and so on.

BASIC **ADVANCED**

General Configuration

Hardware and Storage

3 **Network and Availability**

4 Cloud Storage

5 Cluster Extensions

6 External Sources

7 Gateway Configuration

8 Network Security Groups

9 Security

Select the virtual network in which you would like your cluster to be provisioned. You can define custom network configurations or use default network configurations.

Select Network
vpc-26772d41

Select Subnet
subnet-7811451f

Custom Domain*
mydatalake.local

Custom Hostname*
prod

Previous **Next**

5. On the Cloud Storage page:

- Under Cloud Storage, configure access to cloud storage via the method available for your cloud provider.
- Under Storage Locations, provide an existing location within your cloud storage account that can be used for Ranger audit log. If using HDP 2.6, this will also be used for Hive metastore directory and Hive warehouse directory. If using the HA blueprint, this location will also be used for HBase Root Directory.



Note:

The storage location must exist prior to data lake provisioning. If the storage location does not exist then Ranger is installed properly, but it may not work.

BASIC **ADVANCED**

- General Configuration
- Hardware and Storage
- Network and Availability
- 4 Cloud Storage**
- 5 Cluster Extensions
- 6 External Sources
- 7 Gateway Configuration
- 8 Network Security Groups
- 9 Security

Cloud Storage

You can optionally configure your cluster to access cloud storage data.

Use existing instance profile

Name of an existing instance profile to use for accessing S3 from node instances. The profile must have the s3access policy attached.

123456-domi-s3

arn:aws:iam::315627065446:instance-profile/123456-domi-s3 created on Fri Jun 29 18:31:05 UTC 2018

Storage Locations

Optionally set the following cluster configurations for cloud storage locations.

Configure Storage Locations

Existing Base Storage Location*

s3a:// domi-test

Path for Hive Warehouse Directory property (hive.metastore.warehouse.dir)
Hive stores the table data for managed tables in the Hive warehouse directory

s3a:// domi-test/test-data-lake-ha/apps/hive/warehouse

Path for Ranger Audit Logs property (xasecure.audit.destination.hdfs.dir)
This is the Directory where audit logs should be stored

s3a:// domi-test/test-data-lake-ha/apps/ranger/audit/test-data-lake-ha

6. On the External Sources page, select the previously registered Ranger database, Hive database and LDAP. If using the HA blueprint, also select the previously registered Ambari database:

BASIC **ADVANCED**

- General Configuration
- Hardware and Storage
- Network and Availability
- Cloud Storage
- Cluster Extensions
- 6 External Sources**
- 7 Gateway Configuration
- 8 Network Security Groups
- 9 Security

Configure Authentication

Authentication Configuration List*

test-ldap2 (fwxmsad-bd87e95aa9775a71.elb.eu-west-1.amaz...

For a Data Lake, you must select authentication source.

Configure Databases

Using an external database for a service allows you to preserve the service metadata and reuse it between clusters.

Database Configuration List

There is no unattached database configuration to display.

For Ambari HA, you must attach a database configuration for Ambari.
 For a Data Lake, you must attach a database configuration for Hive and Ranger.

Attached Database Configurations

ranger-db for Ranger hive-db for Hive ambari-rd for Ambari

7. On the Gateway Configuration page, the gateway is enabled by default with Ambari exposed through the gateway. You should also enable Ranger by selecting the Ranger service and clicking Expose.

BASIC **ADVANCED**

- General Configuration
- Hardware and Storage
- Network and Availability
- Cloud Storage
- Cluster Extensions
- External Sources
- 7 Gateway Configuration**
- 8 Network Security Groups
- 9 Security

Gateway Topology

Deploy a protected gateway (Apache Knox) for cluster access.

Name*

dp-proxy

Exposable Services

Select services to be proxied through Knox Gateway.

ALL

Exposed Services

Ambari Ranger

8. On the Network Security Groups page, you do not need to change anything. If you would like to restrict the open ports, refer to [Default cluster security groups](#).
9. On the Security page:
 - Under Password, provide a strong password for your cluster. For example “SomeRandomChars123!” is a strong password. A strong password is required for the default Ranger admin, which - among other cluster components like Ambari - will use this password.
 - Select an SSH key.
10. Click Create Cluster to initiate data lake creation.

As an outcome of this step, you should have a running data lake. Once the data lake is running, you can create workload clusters attached to it.

Related Information

[Creating a cluster on AWS](#)

[Creating a cluster on Azure](#)

[Creating a cluster on GCP](#)

[Creating a cluster on OpenStack](#)

[Default cluster security groups](#)

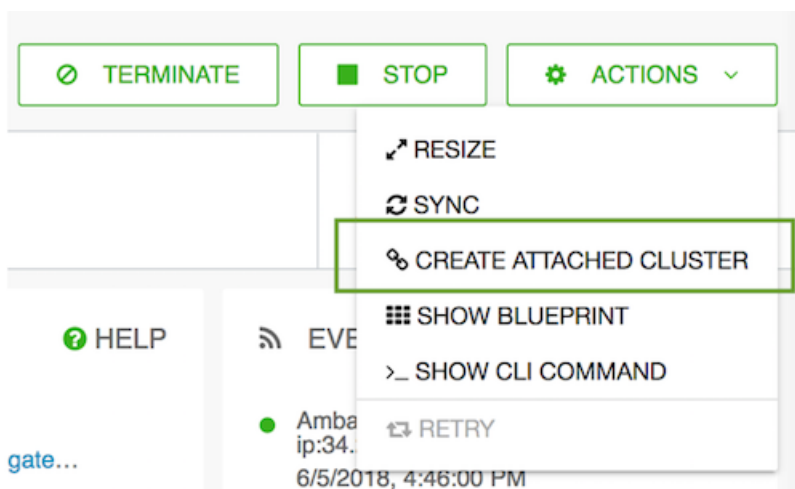
Create attached workloads

Once your data lake is running, you can start creating workload clusters attached to the data lake.

Follow these general steps to create an cluster attached to a data lake. In general, once you've selected the data lake that the cluster should be using, the cluster wizard should provide you with the cluster settings that should be used for the attached cluster.

Steps

1. In the Cloudbreak web UI, click on the cluster tile representing your data lake.
2. From the ACTIONS menu, select CREATE ATTACHED CLUSTER.



3. In general, the cluster wizard should provide you with the cluster settings that should be used for the attached cluster. Still, make sure to do the following:
 - Under Region and Availability Zone, select the same location where your data lake is running.
 - Select one of the three default blueprints.
 - On the Cloud Storage page, enter the same cloud storage location that your data lake is using.
 - On the External Sources page, the LDAP, and Ranger and Hive databases that you attached to the data lake should be attached to your cluster.
 - On the Network page, select the same VPC and subnet where the data lake is running.

4. Click on CREATE CLUSTER to initiate cluster creation.

As an outcome of this step, you should have a running cluster attached to the data lake. Access your attached clusters and run your workloads as normal.

Performing manual HA data lake recovery

When auto repair is not enabled, in case of a node failure, you must perform manual repair.

If on the Hardware and Storage page you did not select to Enable Auto Recovery, you must perform a manual recovery when a data lake host goes down.

Manual repair from web UI


When host-level failures are detected on worker or compute nodes, the following message is displayed on the cluster tile:

```
The cluster has unhealthy nodes
```

In addition, a message similar to the following is written to the EVENT HISTORY, and the status of the node changes from green to red:

```
Manual recovery is needed for the following node...
```

To trigger manual repair, navigate to the cluster details and select Repair from the Actions menu:



Manual repair from CLI

You can perform similar steps in CLI by using the following CLI commands:

- `cb cluster list` – Allows you to check the status and health of your clusters
- `cb cluster describe` – Allows you to check the status and health of a specific cluster
- `cb cluster repair` – Perform cluster repair