Cloudera on Premises Data Services 1.5.5

Cloudera Management Console on on premises Administration

Date published: 2023-12-16
Date modified: 2026-08-26



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2026. 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

Managing Cloudera on premises diagnostic data..... 4
Collecting Cloudera on premises diagnostic data....4
Options for generating the Cloudera on premises diagnostic data....5
Cloudera Data Services on premises diagnostic data....7

Configuring LDAP authentication for Cloudera on premises..... 9

Modifying database properties.....11

Updating TLS certificates.....14
Updating custom certificates....15
Manual configuration for Ranger RAZ certificates....16

Configuring alerts for Cloudera on premises.....16
Configuring alert receivers....16
Configuring alert rules....19
Additional operations on the alert rules....20

NTP proxy setup on Cloudera AI.....21
Installing a non-transparent proxy in a Cloudera AI environment....21
Enabling proxies in Cloudera Embedded Container Service and OpenShift Container Platform environment....23
Proxy setting best practices....25

NTP proxy setup on Cloudera Data Engineering..... 27
Installing a non-transparent proxy in a Cloudera Data Engineering environment....27
Enabling proxies in Cloudera Embedded Container Service and OpenShift Container Platform environment....29
Proxy setting best practices....31

Managing Cloudera on premises diagnostic data

Cloudera on premises enables you to generate and download diagnostic data associated with the various services and workloads for troubleshooting purposes. You can perform these tasks from the Diagnostics tab of the Cloudera Management Console.

Cloudera on premises uses fluentd to collect logs from the Cloudera Management Console services, and the Cloudera AI, Cloudera Data Warehouse, and Cloudera Data Engineering workloads. If Ozone is configured to store logs from the workloads, the information collected by fluentd is stored on Ozone. Otherwise, the data is moved through a flue ntd aggregator in the Cloudera Management Console to an OpenShift Container Platform persistent volume where all the Cloudera Management Console logs are also stored.

When an administrator selects the option to generate the diagnostic bundle from the Cloudera Management Console > Diagnostic page, Cloudera Management Console generates a bundle containing the following details after querying the relevant components and services:

- Logs from the OpenShift Container Platform persistent volume
- Cloudera AI, Cloudera Data Warehouse, and Cloudera Data Engineering logs stored on Ozone
- Kubernetes deployment information such as events, logs, details of the OpenShift Container Platform pods and so on.
- Metrics for monitoring such as usage information and so on.
- Cloudera Management Console information such as version, UUID, and so on.

Collecting Cloudera on premises diagnostic data

You can collect and download Cloudera on premises diagnostic data for different components and services from the Diagnostics page of the Cloudera Management Console. In addition, you can specify the type of workloads and the duration for which you want to collect the diagnostic data.

Procedure

1. Cloudera Management Console home page, select Diagnostics.
   The Diagnostics homepage contains data related to the collected Bundles and Redaction Rules.
2. Under the Diagnostics homepage, you can perform the following actions:

Collect New Bundle and **Add New Rule**

3. Select the Bundles tab and click Collect New Bundle.
   On the Collect New Bundle page, you can specify the following options for data collection: the time duration, the associated environments and services, and the size of the file to include the diagnostic data. For more information on the diagnostic data options, see Options for generating the Cloudera on premises diagnostic data on page 5
4. Click Collect and Send Diagnostic Data to commence the process.

   Cloudera Management Console initiates the process of generating the diagnostic data based on the options that you select. After the data is generated, the Last Collected Time field is updated with the collection time.

   Once the collection process is completed, navigate to the Diagnostics homepage and click on the collected bundle name to fetch more details. Alternatively, you can download or delete the selected bundle. The diagnostic data is downloaded in a .zip file. You can send this file to Cloudera support. If your system is connected to the internet, Cloudera Management Console collects the diagnostic data and sends it to Cloudera Support.
5. Select the Redaction Rules tab and click Add New Rule.

   Redaction Rules define the list of rules to be followed for redacting sensitive information from the diagnostic bundle. Sensitive information, including but not limited to passwords and personal identifiers, is not collected.
6. Input the information for the following:

- Description: Enter the description for the redaction rule. This is a mandatory field.

- Trigger: Enter the redaction rule trigger string.
- Search: Enter the redaction rule search expression. A boolean value that indicates whether the search regular expression should be case-sensitive. This is a mandatory field.
- Replace: Enter the redaction rule replace string. This is a mandatory field.

7. Click Add New Rule.
   The defined redaction rule is created.
8. Select Cancel to return to the previous page.

Options for generating the Cloudera on premises diagnostic data

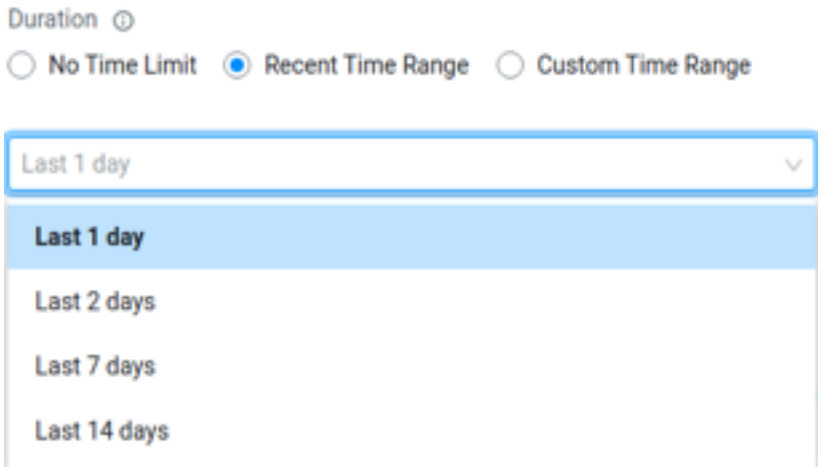
You can specify various options on the Collect and Send Diagnostic Data pop-up window to generate Cloudera on premises diagnostic data based on your specific requirements of duration, scope, and size.

Duration

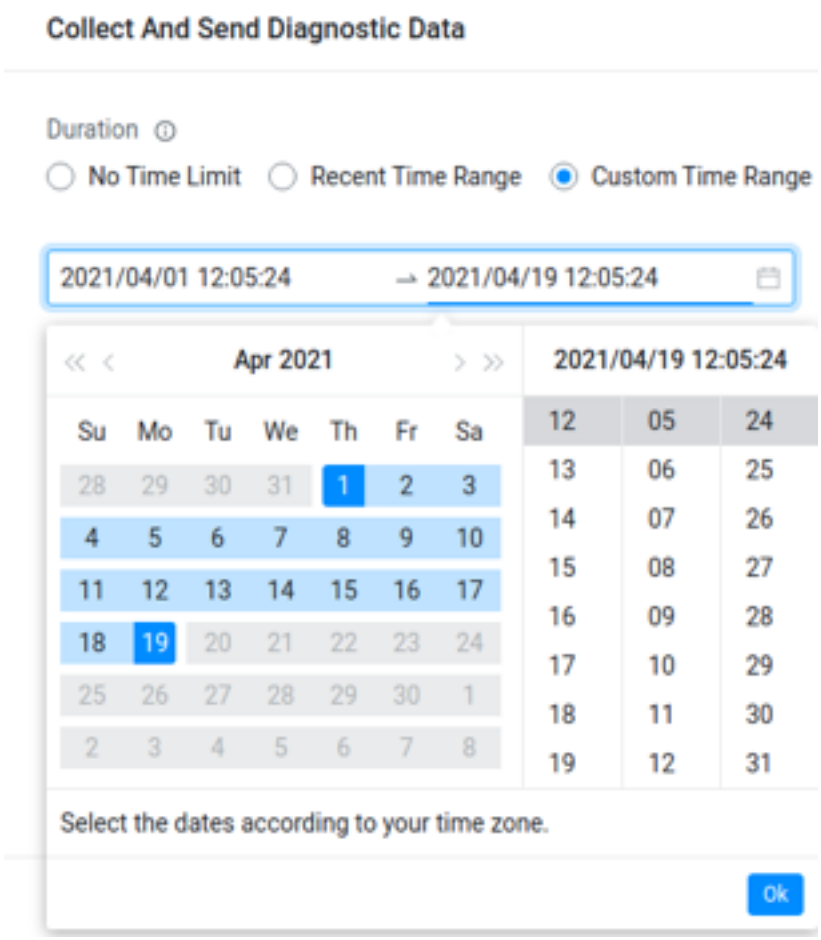
Select the duration for which you want to generate the diagnostic data.

You can select from the following options:

- No Time Limit: Select this value if you want all the diagnostic data to be generated from the time you deployed Cloudera on premises.
- Recent Time Range: Select a specified time range, in hours, from the drop-down list.



- Custom Time Range: Select the start date and end date between which you want to generate the diagnostic data.

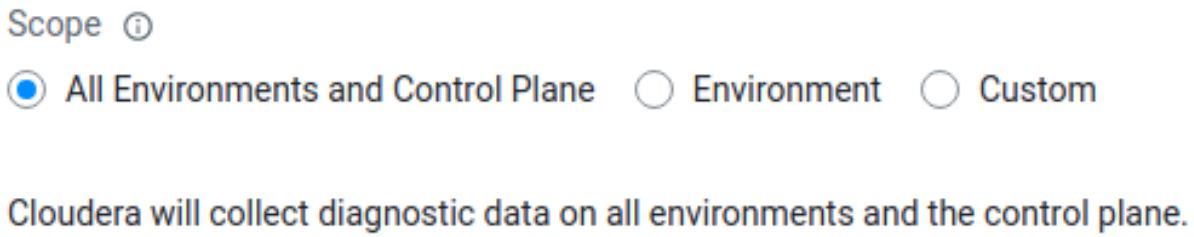


Scope

Select the services for which you want to collect the diagnostic data.

You can select environment- and workload-related options from the following:

- All Environments and Control Plane: Select this option to generate the diagnostic data for all the environments and workloads deployed on Cloudera on premises.



- Environment: Select one or more environments from the drop-down list to generate the diagnostic data. You can also use regular expressions to filter the name of environments for which you want to generate the diagnostic data.
- Custom: You can filter your selections to specific environments and workloads, and accordingly generate the diagnostic data.

  You must click Add Scope to select an environment and provide information about its associated workloads for which you want to generate the diagnostic data. The details include the following:

  - Workload Status: For a selected environment, specify the option to select either a live or an archived workload.

- Workload Type: For a Live workload status, you can specify the option to select Cloudera AI, Cloudera Data Warehouse, or Cloudera Data Engineering.
- Workload Name: For a selected workload type, you can specify the names of the workloads for which you want to generate the diagnostic data. You can add a single workload or multiple workloads as comma-separated values. If required, you can use regular expressions to filter the workload names.
- Namespace: For an Archived workload status, specify the Kubernetes namespace corresponding to the workloads for which you want to generate the diagnostic data. You can add a single namespace or multiple namespaces as comma-separated values. If required, you can use regular expressions to filter the namespace values.
- Pod: For an Archived workload status, specify the OpenShift Container Platform pods on which the workloads are archived. You can add a single pod or multiple pods as comma-separated values. If required, you can use regular expressions to filter the names of the pods.



Note:

- If you do not specify any namespace, workload name, or pod, Cloudera Management Console generates diagnostic data for all the namespaces, workloads, and pods associated with the specified environment-workload type combination.
- Ensure that any regular expression that you use on the Collect and Send Diagnostic Data pop-up window is anchored by a ^ symbol at the beginning of the expression and a $ symbol at the end.

Storage

You can specify the size of the diagnostic data to collect. The default size of the data after compression is 500 MB. If you want to change the default size, set the BUNDLE_SIZE_LIMIT_MB environment variable on the OpenShift Container Platform pods on your Cloudera on premises deployment.



Note: The actual size of the generated diagnostic bundle might vary from the size that you specify.

While Cloudera can exclude namespace pod logs, control plane archive logs, and other archive logs from the diagnostic bundle to be consistent with the specified size, Cloudera must always include certain type of diagnostic data in the bundle. Because the size of the diagnostic data that is always included in a bundle is not constant, the overall size of the generated diagnostic bundle might vary.

Cloudera Data Services on premises diagnostic data

You can collect and download Cloudera Data Services on premises diagnostic data using the Bundles tab on the Diagnostics page of the Cloudera Management Console. This topic lists the diagnostic data that is generated.

Cloudera Data Services on premises diagnostics bundle

Namespace information (under <namespace>/):

- pods.json
- events.json
- replication.json
- services.json
- daemonset.json
- deployments.json
- replicasets.json
- configmap.json
- pv.json
- Pod logs (under logs/)
- Archive logs (under archived/)

bundle.json

- File name of diagnostics bundle
- Control plane version
- Control plane namespace

Environment information (under environment/)

- Per environment (under <environment-name>/)
  - Archive Logs (under archived/)
    - Non-Ozone logs (those stored in the persistent volume claim of the fluentd-aggregator pod)
    - Ozone logs (those stored in Ozone)
- cde
  - Separated into folders per namespace for the Cloudera Data Engineering services
  - Each folder contains namespace information as above
- dwx
  - Separated into folders per Cloudera Data Warehouse containing namespace information
  - DWX diagnostics bundle for each warehouse
- mlx
  - Separated into folders per MLX workspace
  - Each folder contains namespace information as above
- environment.json
  - Contains the following:
    - CM version,
    - CDH version
    - CM License UUID
    - Health of all of the hosts
- client-configs.zip
  - Client configs of the base cluster
- details.json
  - Environment details

license.json

- License version
- License Name
- License UUID
- License Start date
- License Deactivation date
- License Expiration date

log.txt

- Good for debugging the diagnostics collection
- All log statements from the diagnostics collection are found here

monitoring-metrics.json

- The results of various monitoring queries

nodes.json

- Details of all cluster nodes
- Examples: labels, creationTimeStamp, etc.

Pod reaper logs (under pod-reaper/)

- Job logs for pod reaper

Configuring LDAP authentication for Cloudera on premises

You can configure LDAP user authentication for Cloudera on premises from the Administration page of the Cloudera Management Console.

Before you begin

If you intend to use Hue as your SQL editor in Cloudera Data Warehouse, you must use LDAP over SSL.

Procedure

1. Sign in to the Cloudera console.
2. Click Cloudera Management Console.
3. On the Cloudera Management Console home page, select Administration>Authentication.
4. Configure the following settings for LDAP authentication:

Property	Description	Sample values
LDAP URL	The LDAP server URL. The URL must be prefixed with ldap:// or ldaps://. The URL can optionally specify a custom port, for example: ldaps://ldap_server.example.com:1636. **Note:** Cloudera recommends that you use Active Directory (AD) Global Catalog ports 3268 and 3269 if you are using LDAP referrals. **Note:** If you do not use Global Catalog port, environment activation fails with "Unable to create cluster initial state: Active Directory servers should be used through the Global Catalog ports 3268/3269" error.	ldap://<ldap-host>:389 or ldaps://<ldap-host>:636 For Active Directory use: `ldap://<AD Server>:3268 or ldaps://<AD Server>:3269`

CA Certificate for Secure LDAP	The X.509 PEM certificate to be used to access secure LDAP (URLs starting with ldaps://). Ensure that at least one valid certificate is provided. A typical CA certificate is structured as follows:
	
	`-----BEGIN CERTIFICATE-----`
	`...`
	`-----END CERTIFICATE-----`
		If you add or update CA certificates and you have deployed the Cloudera Data Warehouse service in your ECS cluster, you must refresh the affected Database Catalogs and Virtual Warehouses from the Cloudera Data Warehouse. Go to the Cloudera Data Warehouse UI, click on the more vertical menu on the Database Catalog or Virtual Warehouse and click Refresh.

LDAP Bind DN	The Distinguished Name of the user to bind to LDAP for user authentication search/bind and group lookup for role authorization.	Distinguished Name (DN) example: CN=cdh admin,OU=svcaccount,DC=example,DC=com FreeIPA example: uid=username,cn=users,cn=accounts,dc=example,dc=com
LDAP Bind Password	The bind user password.	

Property	Description	Sample values
LDAP User Search Base	The distinguished name indicating the path within the directory information tree from which to begin user searches.	AD example: cn=users,dc=example,dc=com LDAP example: ou=people,dc=example,dc=com FreeIPA example: cn=accounts,dc=example,dc=com
LDAP User Search Filter	The search filter to use for finding users.	AD example: If you want to filter users belonging to specific group say MLgroup, then set the filter as: "" (&(sAMAccountName={0}) (memberof=CN=MLgroup,OU=Users,OU=XX,OU=XXX,DC= """" LDAP example: (uid={0}) Note that a custom attribute can also be used if the directory is configured differently for user names. The {0} expands the currently authenticating user's name entered in the login form for the query. FreeIPA example: (&(uid={0})(objectClass=person))
LDAP Group Search Base	The distinguished name indicating the path within the directory information tree to begin group searches from.	cn=accounts,dc=example,dc=com

LDAP Group Search Filter	The search filter to use for finding groups for authorization of authenticated users for their roles. You must configure this value such that only the groups associated with the user logging in are fetched from the IdP.
	There are two placeholders that can be used to match the groups of a user, {0) and {1}. {0} expands into the user DN and {1} expands into the username.
		For Active Directory and openLDAP compatible directories this will usually be (member={0}), where {0} will be replaced by DN string for a successfully authenticated user through the search/bind process. This requires configuration of the LDAP Bind User Distinguished Name field.
		AD example:
		`(member={0})`
		LDAP/FreeIPA example:
		`(&(member={0})(objec tClass=posixgroup)(! (cn=admins)))`

Email Mapping Attribute	The LDAP attribute to be used for mapping the email in Identity Management. If no value is provided, mail is used as the default email mapping attribute. Email is a mandatory value in Cloudera. If no value is found for the email attribute, a value of {userame}@cdp.example is assumed for the user.	

Property	Description	Sample values
Username Mapping Attribute	The LDAP attribute to be used for mapping the userId in Identity Management. **Important:** This property must be provided.	The Username Mapping Attribute must be configured as sAMAccountName. An incorrect value (for example, (sAMAccountName={0})) may result in duplicate users being created in UMS, which could cause user access and authorization issues.

Note: For information on setting search filters for nested groups, see *Nested LDAP Group Resolution*.

5. If required, select Show Other Options and configure the following additional settings:

Property	Description	Sample values
LDAP User Bind Property	The property of the LDAP user object to use when binding to verify the password.	This value should always be set to dn.
Groupname Mapping Attribute	The LDAP attribute to be used for mapping the groupId in Identity Management.	
Group DN Property	The property of user object to use in {{dn}} interpolation of groupSearchFilter.	This value should always be set to dn.
First Name Mapping Attribute	The LDAP attribute to be used for mapping the first name attribute in Identity Management.	
Last Name Mapping Attribute	The LDAP attribute to be used for mapping the last name attribute in Identity Management.	

Note: If the username and email attributes on the LDAP server contain spaces, Cloudera Management Console includes the spaces when copying the corresponding attribute information.

6. Click Test Connection to verify whether the LDAP information you have provided is valid.

   Cloudera Management Console attempts a connection with the LDAP source based on the information provided, and returns a confirmation message if the connection is successful.
7. Click Save. The LDAP users are listed on the User Management page.

Important: In Cloudera Data Services on premises 1.5.5 SP3 and higher releases, if you update LDAP configurations after creating Cloudera Data Engineering Service or Virtual Cluster, you must update and synchronize these changes in the Cloudera Data Engineering Service. For more information, see Updating and synchronizing stale configurations.

Related Information

Nested LDAP Group Resolution

Modifying database properties

You can modify external database properties such as hostname, port, username, password, and database name using the Cloudera Management Console.

About this task

If your cluster was installed using the OpenShift Container Platform Platform, you can view, but not modify the database properties. If the cluster was installed using the Cloudera Embedded Container Service, you can edit the properties.

Important:

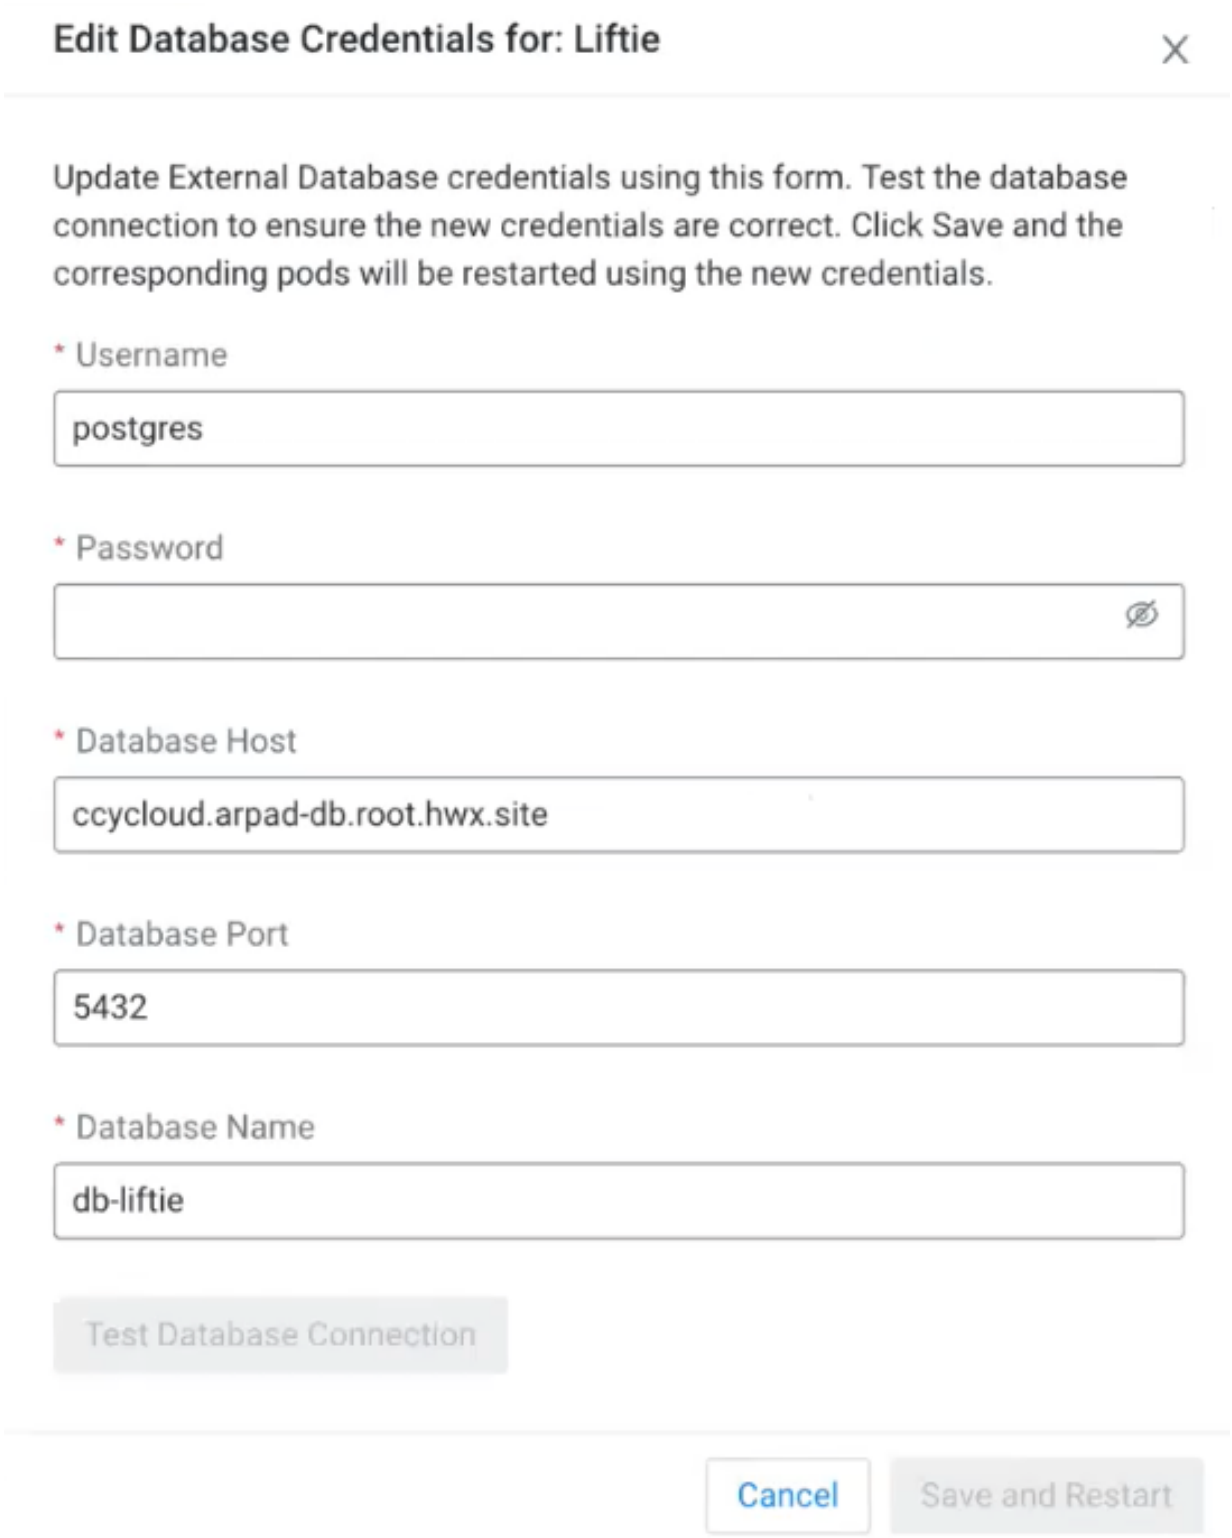
When updating the database password, if the User Management Service (UMS) and other Cloudera Data Services on premises share the same database, then you must update the UMS database password first, followed by other services that share the database.

Procedure

1. Sign in to the Cloudera console.
2. Click Cloudera Management Console.
3. On the Cloudera Management Console home page, select Administration Databases to view the **Database administration** page.

   The **Database administration** page displays a list of databases.
4. Click  located at the right side of each row, for the database you want to modify.

   A dialog box appears where you can modify database properties. For example:



5. Modify any of the following properties:
   - Username
   - Password
   - Database Host
   - Database Port
   - Database Name

6. Click the Test Database Connection button to verify that Cloudera can connect to the database. If the test fails, re-check your modifications and click Test Database Connection again.
7. Click Save and Restart.
8. If you have deployed the Cloudera Data Warehouse service in your Cloudera Embedded Container Service cluster, you must refresh the affected Database Catalogs and Virtual Warehouses from the Cloudera Data Warehouse UI. Go to the Cloudera Data Warehouse UI, click ⋮ on the Database Catalog or Virtual Warehouse and click Refresh.

Results
The database properties are updated.

Updating TLS certificates

From the Cloudera Management Console of your Cloudera on premises deployment, you can update the CA certificate that issued the TLS certificates used by Cloudera to establish secure connections with different types of services and workloads, such as the ingress controller certificate. For specified services, you can update the certificates whenever you rotate them.

Before you begin
You must ensure that all the services for which you want to update the certificates are TLS-enabled.

About this task

Important: The following procedure causes a restart of the dependent services after you have added the certificate. Therefore, Cloudera recommends that you plan for a down time before performing this procedure.

Procedure

1. Sign in to the Cloudera console.
2. Click Cloudera Management Console.
3. On the Cloudera Management Console home page, select Administration>CA Certificates.
4. From the CA Certificate Type dropdown list, select the type of service for which you want to upload a new TLS certificate.

   You can select from the following options for the types of secure connections:

   - Datalake: For secure connections with the Cloudera Base on premises cluster services and Cloudera Manager.
   - Docker Registry: For a secure connection with the Docker Container registry containing the images for deployment.
   - External Database: For a secure connection with an external PostgreSQL database.
   - External Vault: For a secure connection with an external vault.
   - Miscellaneous: For a secure connection with services used during the installation and run time of Cloudera. For example, Custom Ingress, Custom Kubernetes API, and so on.

     Important: If your Cloudera on premises deployment uses an external vault, then after updating the certificates for these services, you must reconfigure the certificates with the vault to ensure that it validates the certificates. For more information, see Updating custom certificates on page 15
5. Select the option to either browse and upload a certificate or directly enter the certificate details.

   Important: The certificate must be in the X.509 PEM format and structured as follows:

   ----BEGIN CERTIFICATE-----..-----END CERTIFICATE----- .

   If you are using AutoTLS, after rotating the certificates, the CA certificate you need to add is located on the Cloudera Manager server at /var/lib/cloudera-scm-agent/agent-cert/cm-auto-in_cluster_ca_cert.pem.

Updating custom certificates

In a Cloudera Private Cloud deployment configured with an external vault, the Cloudera Management services such as Custom Ingress and Custom Kubernetes API authenticate to the vault with the help of the ServiceAccount's JSON Web Token (JWT) attached to the Kubernetes pod on which the services are running.

The vault service validates the JWTs using the Kubernetes TokenReview API, and verifies the existence of the services. However, if the external service dependency being updated uses an entirely new CA certificate that Cloudera Data Services on premises is not currently configured to trust, then that CA certificate should be updated in Cloudera Data Services on premises first, then the certificate of the service dependency can be updated

Before you begin

You must have updated the certificates for the services by using the Miscellaneous certificate type option mentioned in Updating TLS certificates on page 14.

Procedure

- Enter the following cURL statement to authenticate to the vault.

```
$ curl \
      --header "X-Vault-Token: <VAULT_TOKEN>" \
      --request POST \
      --data @payload.json \
      http://<VAULT_URL>/v1/auth/<KUBERNETES_PATH>/config
```

Here:

- <VAULT_TOKEN>: The privileged authorization token with write permissions on the vault.
- <KUBERNETES_PATH>: The path on which the Kubernetes login credentials information is mounted on the vault service. You can find this information specified as the value of the VAULT_AUTH_PATH property in the vault configmap. Ensure that this value follows the following naming convention: cloudera-<PROJECT_NAME>-<K8s-host> with the dot (.) replaced by an underscore (_).
- <VAULT_URL>: The URL of the vault service.

The payload.json contains the following information:

- kubernetes_host: The URL to access the Kubernetes API server from the vault service.
- token_reviewer_jwt: The JWT of the Kubernetes service account that the vault service uses to validate authentication requests from the Management Console services. The Cloudera Private Cloud installer creates a dedicated service named vault-auth for reviewing the requests.
- kubernetes_ca_cert: The CA certificate of the Kubernetes API server with newline characters replaced with '\n'.

The following example shows the contents of payload.json:

```
{
  "kubernetes_host": "https://api.examplehost.com:1111",
  "token_reviewer_jwt": "-----BEGIN CERTIFICATE-----\n.......\n-----END C
ERTIFICATE-----",
  "kubernetes_ca_cert": "-----BEGIN CERTIFICATE-----\n.......\n-----END CE
RTIFICATE-----"
}
```

For more information, see the vault documentation.

Manual configuration for Ranger RAZ certificates

When configuring custom TLS certificates on a Cloudera on premises deployment with Ranger RAZ enabled, administrators must manually upload the Ranger RAZ CA certificate

About this task

By design, a Virtual Warehouse within Cloudera Data Warehouse fetches its trusted certificates from the Cloudera Management Console during provisioning. The Ranger RAZ CA certificate is not included by default due to strict service isolation boundaries that prevent automated certificate propagation. Leaving this unconfigured will cause secure handshakes to fail, blocking data operations like INSERT queries against external tables stored at S3.

To ensure proper cross-service communication, use the steps below to fetch the Ranger RAZ root certificate and import it as a miscellaneous certificate type.

Procedure

1. Run the following general openssl command from a terminal that has network access to the Ranger RAZ host to extract its certificate:

```
echo | openssl s_client -showcerts -connect <RAZ_HOST>:<RAZ_PORT> 2>/dev
/null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > <PATH_TO_
SAVED_CA_CERT>
```

```
echo | openssl s_client -showcerts -connect kg-rzmtr-1.kg-rzmtr.root.hwx
.site:6082 2>/dev/null | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICAT
E-/p' > /tmp/ranger_raz_ca.crt
```

2. Add the certificate to the Cloudera Management Console:
   a) In the Cloudera Management Console, navigate to Administration CA Certificates .
   b) Click Add/Import.
   c) Select the CA Certificate Type as Miscellaneous.
   d) Click Choose file to upload the fetched certificate.
   e) Click Add certificate.

Results

The Ranger RAZ certificate is displayed in the CA Certificates tab.

Configuring alerts for Cloudera on premises

Cloudera Management Console helps you configure alert receivers and rules to trigger automated notifications when specific events occur in your Cloudera on premises deployment. You can view these alert notifications on the Cloudera Management Console dashboard. Further, you can send the notifications to specified endpoints.

Configuring alert receivers

You can configure alert receivers on Cloudera Management Console to trigger automated system-specific event notifications through external services such as emails, Slack channel messages, webhook notifications, PagerDuty messages, or SNMP traps.

By configuring an alert receiver, you specify the details of an external service through which Cloudera Management Console forwards the notification to the specified destination.

Before you begin

- You can configure alert receivers *only* if you have administrator privileges.
- Make sure you have the details of the external services for which you are configuring the alert receiver.
- If your Cloudera on premises deployment is air gapped, you must have services available within your internal network that can help forward the notifications. For example, you must have one or more of the following services internally available: an SMTP server, a gateway or a proxy that can forward the notifications to external systems, a webhook endpoint that can externally forward the notifications, or an SNMP server.

  Further, you must configure your firewall policies to enable outgoing traffic to external systems.

Procedure

1. Sign in to the Cloudera console.
2. Click Cloudera Management Console.
3. On the Cloudera Management Console home page, select Administration>Alerts to view the Alerts page.
4. In the Alert Receivers section of the page, click Add Alert Receiver.
5. Configure the following options for adding an alert receiver in the pop-up window:



Note:

- You can configure *only one* alert receiver of type SNMP at a time.
- If your Cloudera Data Services on premises cluster is deployed using the Cloudera Embedded Container Service (ECS), then Cloudera Manager automatically imports at the time of installation any non-default SMTP and SNMP alert configurations from the Cloudera Base on premises cluster. Accordingly, alert receivers of type SNMP and SMTP are automatically configured when deploying Cloudera Management Console for the first time.

Field	Description
Alert Scope	
Severity	The severity of the alert notifications to configure. You can select notifications of either Warning, or Critical, or both the types.
Source	The Cloudera Management Console services for which you want to configure the alert notifications. You can select to configure notifications for the environments, or the control plane, or all the environments and control plane.
Receiver Type	The type of external service that receives the alert notification from Cloudera Management Console. You can select from the following types: • Email • Slack • Webhook • PagerDuty • SNMP
Email	
Email To	The destination email address to which the alert notification email must be forwarded.
Email From	The source email address from where the alert notification email must be forwarded.
SMTP Server	The URL of the SMTP server through which the alert notification email is forwarded. You must specify the host name and port number as part of the SMTP server URL.
User name	The username to access the SMTP server.
Password	The password to access the SMTP server.

Field	Description
Connection Requires TLS	Toggle this option to On if you have a secure connection configured between your Cloudera on premises deployment and the SMTP server such that the SMTP server's certificate is trusted by Cloudera on premises.
Slack	
API URL	The URL of the Slack API associated with the channel to which the notification must be sent.
Channel	The name of the Slack channel. You can specify any name here for your reference because the Slack API derives the actual name from the specified URL.
Webhook	
URL	The URL of the webhook endpoint to which the alert notification must be sent.
Basic Auth User Name	The username for basic authentication to the webhook application.
Basic Auth Password	The password for basic authentication to the webhook application.
Bearer Token	The authentication header in case of using a bearer token.
Skip TLS Verification (Insecure)	Toggle this option to On if you have a secure connection configured between your Cloudera on premises deployment and the webhook application. For the secure connection, you must use a certificate trusted by Cloudera on premises.
PagerDuty	
URL	The URL of the PagerDuty platform to which the notification must be sent.
Routing Key	The PagerDuty Events API v2 integration key. For details, see PagerDuty Services and Integrations.
SNMP	
SNMP NMS Hostname	The DNS name or IP address of the SNMP Network Management Software (NMS) host listening for SNMP traps or notifications.
SNMP Server Port	The port number on which the SNMP server is listening for traps or notifications.
SNMP Retry Count	The maximum number of times to try an SNMP trap before the latter times out. If you specify '0', the trap is sent only once.
SNMP Timeout (Milliseconds)	The time, in milliseconds, to wait after which an SNMP trap times out.
SNMP Security Level	The level of security to use if you the select SNMPv3 protocol for the alert receiver. You can select either authNoPriv or noauthNoPriv security level. The details that you need to specify vary with the security level that you select. You can also select the SNMPv2 protocol, if required.
SNMPv2 Community String	The community string for identifying the SNMP service. Generated SNMPv2 traps use this string for authentication purposes. Specify this value if you select SNMPv2 as the SNMP Security Level.
SNMP Authentication Protocol	The authentication algorithm. The available options are MD5 and SHA. Specify this value if you select authNoPriv as the SNMP Security Level.
SNMP Server Engine Id	Used along with the pass phrase to generate keys for authentication and privacy protocols. The Engine ID is a hexadecimal number. Specify this value if you select authNoPriv as the SNMP Security Level.

Field	Description
SNMP Security Username	The name of the user to add for SNMP security. Specify this value if you select either authNoPriv or noauthNoPriv as the SNMP Security Level.
SNMP Authentication Protocol Pass Phrase	The pass phrase to use for the SNMP authentication protocol. Specify this value if you select authNoPriv as the SNMP Security Level.
Test receiver parameters	
Send Test Event	After providing the details to configure the external service, you can send a test notification to verify whether the message reaches the desired destination.

6. Click Add Alert Receiver.

   The Alert Receivers page now displays the details of the receiver that you just added.
7. Repeat steps 4 and 5 to add more alert receivers.



Note: If you want to edit the details of an alert receiver, select Edit Alert Receiver from the drop-down menu against that receiver entry, and update the desired fields. You can access the drop-down menu by clicking the vertical ellipsis (three dots) against the particular receiver entry on the table.

Similarly, if you want to delete an alert receiver, select Delete Alert Receiver, from the drop-down menu against that receiver entry.

Configuring alert rules

You can define alert rules for your Cloudera on premises deployment based on PromQL expressions. The alerts are automatically triggered when specific events occur in your deployment. You can view the triggered alerts on the Cloudera Management Console dashboard. Any alert receivers that you have already configured start sending notifications to specified endpoints.

Before you begin

You can configure alert rules *only* if you have administrator privileges.

About this task

Cloudera Management Console supports two types of alert rules: built-in and custom.

- Built-in alert rules are system-generated, and therefore, you cannot add or remove them. You can only enable, disable, or edit certain details associated with them.
- You can create custom alert rules based on your requirements. Further, you can edit, delete, enable, or disable them.

Procedure

1. To access the Cloudera Management Console, click Data Services in Cloudera Manager, then click Open Cloudera on premises , then click Cloudera Management Console.
2. On the Cloudera Management Console home page, select Administration > Alerts.
3. On the Alerts page, click Add Alert Rule.
4. Configure the following options for adding an alert rule in the pop-up window:

Field	Description
Name	The name of the alert rule. You cannot use spaces or special characters in the name.
Severity	Specify the severity: Critical or Warning.
Enable Alert	Select this checkbox to enable the alert rule.

Field	Description
Message	The alert rule text. You can use PromQL labels to denote entities such as jobs in the text. For more information about using PromQL labels, see Alerting Rules.
Summary	A brief summary of the alert rule.
Time Period	The duration for which the PromQL expression must be true. If the expression continues to be true after the specified duration, the configured alert is automatically triggered.
Source	The Cloudera on premises components for which the alert rule applies. You can select one of the following options as the source: • All Environments and Control Plane • Environments • Control Plane • A specific environment from the list of configured environments
PromQL Expression	The query expression in PromQL. The alert is issued when this expression is true for the time period specified in the For Clause. **Important:** Metrics reported by the environments always contain the following labels: appId and appName. Therefore, the result of the alert rule's query expression also must contain these labels. To ensure that the result contains the labels, include the by (appId, appName) clause when using aggregation operators in the query expression. For example, instead of the count(my_metric) > 0 expression, use the count(my_metric) by (appId, appName) > 0 expression.
Generated Query	The query that is generated for a selected workload type depending on the specified PromQL expression. You can view the query by clicking Show Generated Query.
Test PromQL Expression	You can click this option to test the query expression generated for the combination of a selected source and workload type. If you select one of All Environments and Control Plane, Environments, or Control Plane as the source for the PromQL query, it runs *only* on the control plane. To run the query on an environment, you must select a specific environment as the source. **Note:** If you click this option *before* saving the alert rule and if the PromQL expression is invalid, an unexpected error appears. In addition, you may lose all of the information entered for configuring the alert rule. Therefore, you should save the alert rule and then test the PromQL expression.

5. Click Add Alert Rule.

   The new alert rule is listed on the Alerts page under Alerting Rules.
6. Repeat steps 4 and 5 to add more alert rules.

Additional operations on the alert rules

You can perform different operations on the added alert rules from the Alerts page.

- If you want to edit the details of an alert rule, select Edit Alert Rule from the drop-down menu against that rule entry, and update the desired fields. You can access the drop-down menu by clicking the vertical ellipsis (three dots) against the particular rule entry on the table.

The fields that you can edit vary with the type of the alert rule. For example, you can edit only the Severity, For Clause fields, and select or clear the Enable Alert checkbox for a built-in alert rule. However, you can edit almost all the available fields for a custom alert rule.

- If you want to enable an alert rule or disable an already enabled rule without opening the pop-up window, select Enable Alert Rule or Disable Alert Rule from the drop-down menu against that rule depending on your requirement.
- If you want to delete a custom alert rule, select Delete Alert Rule from the drop-down menu against that rule.

NTP proxy setup on Cloudera AI

Cloudera AI requires specific proxy configurations to manage workbench connections efficiently in an air-gapped setup with restricted outbound connections. This setup ensures seamless access to external resources while adhering to network security and management policies.

Depending on your cluster platform, whether it is RKE2 for Cloudera Embedded Container Service clusters or OpenShift Container Platform, specific configurations and deployment methods apply. Consider the detailed instructions for configuring proxy settings in Cloudera Embedded Container Service server and agent configurations, as well as enabling cluster-wide proxies in OpenShift Container Platform environments. Also consider the provided sample configurations and guidelines for configuring proxy servers, including specifying Classless Inter-domain Routing (CIDR) ranges to exclude from proxy routing and updating proxy server allowlists.

Installing a non-transparent proxy in a Cloudera AI environment

If Cloudera AI is used in an air-gapped environment, a proxy configuration is not mandatory. If a non-transparent proxy is used, then certain endpoints must be added to the list of allowed endpoints for the proxy.

Before you begin

Configure the No Proxy value with the Classless Inter-Domain Routing (CIDR) ranges for the Nodes, POD CIDR, and Service CIDR. Any IP range for internal services with seamless internal network connectivity must be added in the No Proxy configuration. Specify these CIDR ranges in the configuration to ensure that the traffic destined for these ranges bypasses the proxy. Add comma-separated no-proxy configurations without any spaces between them.

About this task

If your Cloudera on premises deployment uses a non-transparent network proxy, configure proxy hosts that the workloads can use for connections with Cloudera AI Workbenches. You can configure the proxy configuration values from the Cloudera Management Console.

The procedure for updating these settings might be different and dependent on the proxy server software used.



Important:

Apply these procedures before installing or upgrading the Cloudera AI Workbenches.



Note: The configured settings apply for Cloudera on premises deployment using the Cloudera Embedded Container Service. In an OpenShift Container Platform deployment, the default values are used.

Procedure

1. Sign in to the Cloudera console.
2. Click Cloudera Management Console.
3. On the Cloudera Management Console home page, select Administration Networks to view the **Networks** page.
4. Configure the following options for the proxy values:

Table 1: Proxy values

Field	Description
HTTPS proxy	It is the HTTP or HTTPS proxy connection string used with the CML workspaces. You must specify this connection string in the form: http(s)://[***USERNAME***]:[***PASSWORD***]@[***HOST***]:[***PORT***]. The [***USERNAME***] and [***PASSWORD***] parameters are optional. You can specify the connection proxy string without these parameters.
HTTP proxy	It is the HTTP or HTTPS proxy connection string used with the Cloudera AI Workbenches. You must specify this connection string in the form: http(s)://[***USERNAME***]:[***PASSWORD***]@[***HOST***]:[***PORT***]. The [***USERNAME***] and [***PASSWORD***] parameters are optional. You can specify the connection proxy string without these parameters.
No proxy	This is a comma-separated list of hostnames, IP addresses, or hostnames and IP addresses that should not be accessed through the specified HTTPS or HTTP proxy URLs. In case of Cloudera Embedded Container Service deployments, you must include no-proxy URLs for the following: • All the Cloudera Embedded Container Service hosts in your deployment • Any Cloudera Private Cloud Base cluster that you want to access • CIDR IP addresses for internal operations in the Cloudera Embedded Container Service cluster: 10.42.0.0/16 and 10.43.0.0/16 • The .apps domain in the following format: .apps.[***APPS DOMAIN NAME***].

5. Click Save.
6. Ensure that the following endpoint is allowed:

Table 2: Endpoint details

Description	CDP Service	Destination	Protocol and authentication	IP protocol/ Port	Comments
Accelerators for ML Projects (AMPs)	Cloudera AI	https://raw.githubusercontent.com https://github.com	HTTPS	TCP/443	Files for AMPs are hosted on GitHub.

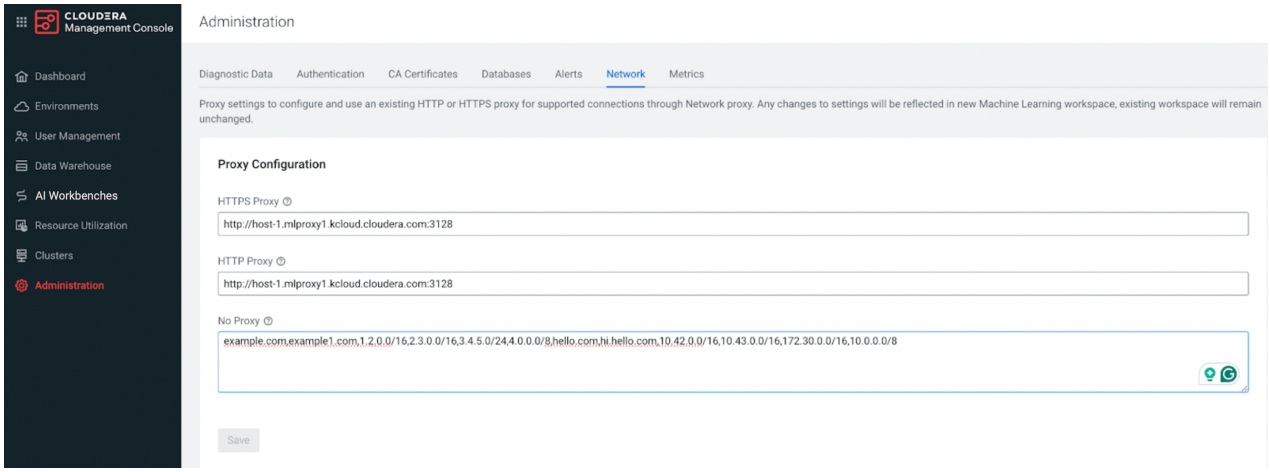
Additionally, ensure that the proxy server's allowlist includes the following specific URLs, which requires updates to the proxy server configuration:

- Cloudera AI Workbench URL, for example: ml-samplexxxx.host-1.proxy.kcloud.cloudera.com
- Cloudera console URL, for example consoles.ml-samplexxxx.apps.host-1.proxy.kcloud.cloudera.com
- External registry if used in Cloudera AI

Example

Consider the following example:

Figure 1: NTP Proxy configuration example



Enabling proxies in Cloudera Embedded Container Service and OpenShift Container Platform environment

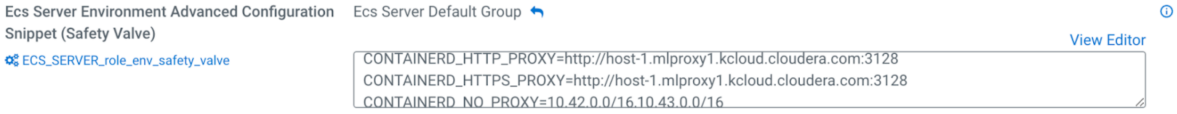
Depending on your cluster platform, whether it is RKE2 for Cloudera Embedded Container Service clusters or OpenShift Container Platform, specific configurations and deployment methods apply.

For ECS advanced configuration settings

1. Open Cloudera Manager Clusters Configuration Configuration Search .
2. Enter the name of the configuration you are looking for: ECS Server Environment Advanced Configuration.

   With the help of the Settings menu, include the environment variables in uppercase in the ECS Server Envir onment Advanced Configuration snippet (Safety Valve) for config.yaml to configure the CONTAINERD_H TTP_PROXY, CONTAINERD_HTTPS_PROXY and CONTAINERD_NO_PROXY.

   Figure 2: Environment variable settings in ECS Server Environment Advanced Configuration

   Example:

   CONTAINERD_HTTP_PROXY=http://proxy.example.com:8080

   CONTAINERD_HTTPS_PROXY=http://proxy.example.com:8080

   CONTAINERD_NO_PROXY=noproxy.address.com

   Note: Avoid using colons, double quotes, or white spaces as delimiters.

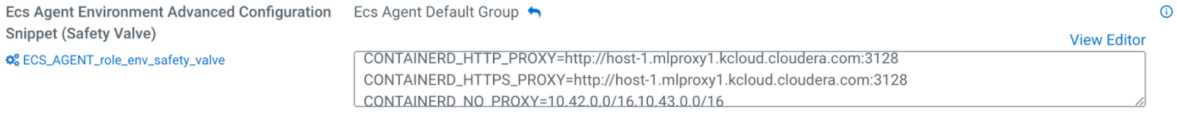
   Note:

   Ensure that CONTAINERD_NO_PROXY contains all the same hosts as the Cloudera on premises configuration.

   Include the host that holds the docker registry that serves all the normal Cloudera Embedded Container Service bundled images.
3. Enter the name of the configuration you are looking for: ECS Agent Environment Advanced Configuration.

With the help of the Settings menu, include uppercase environment variables in the ECS Agent Environment Advanced Configuration snippet (Safety Valve) for config.yaml to configure the CONTAINERD_HTTP_PROXY, CONTAINERD_HTTPS_PROXY and CONTAINERD_NO_PROXY.

Figure 3: Environment variable settings in ECS Agent Environment Advanced Configuration



Example:

CONTAINERD_HTTP_PROXY=http://proxy.example.com:8080

CONTAINERD_HTTPS_PROXY=http://proxy.example.com:8080

CONTAINERD_NO_PROXY=noproxy.address.com

Note: Avoid using colons, double quotes, or white spaces as delimiters.

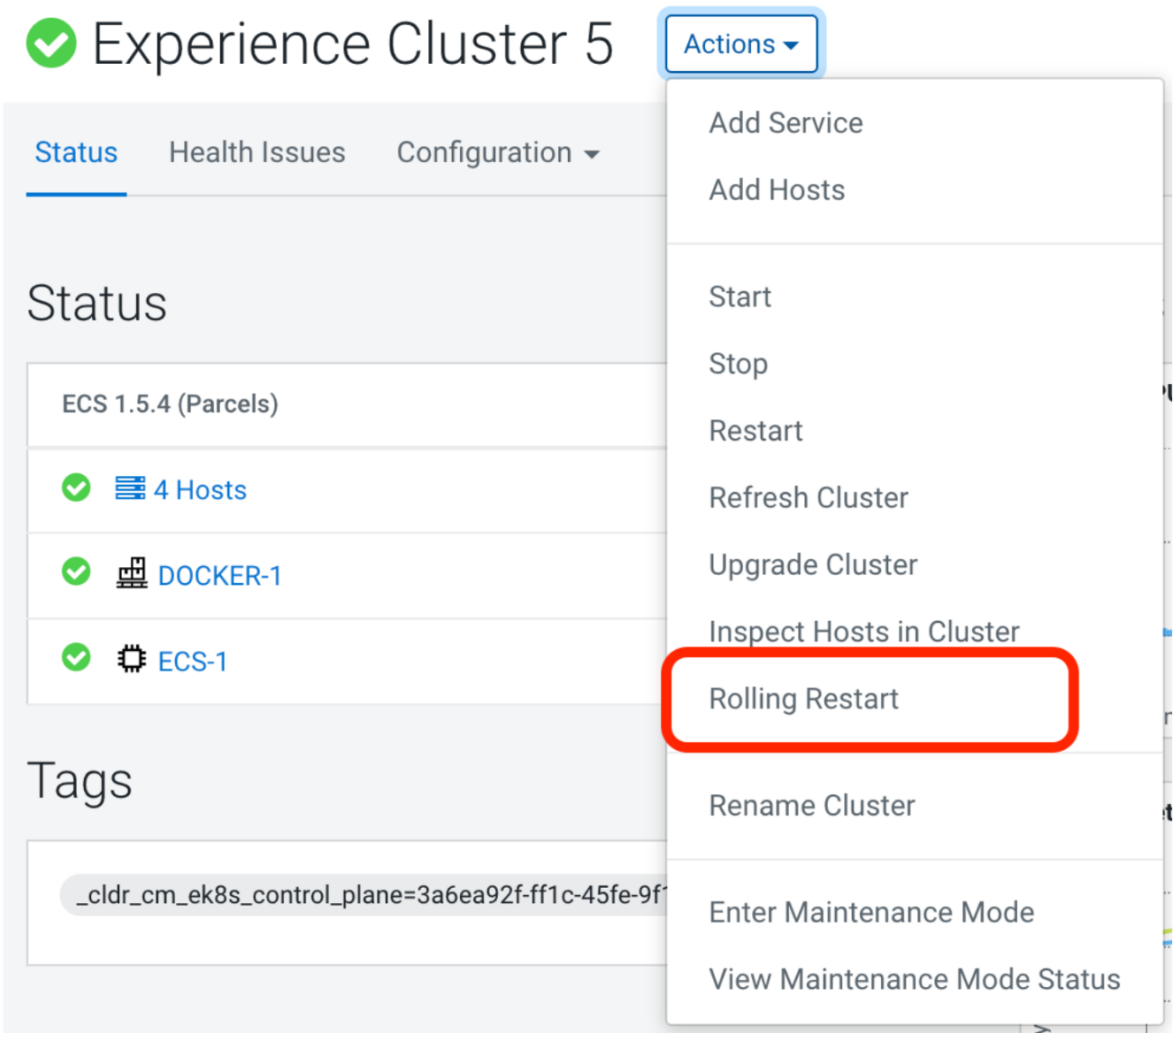
Note:

Ensure that CONTAINERD_NO_PROXY contains all the same hosts as the Cloudera on premises configuration.

Pay special attention to include the host that holds the docker registry that serves all the normal Cloudera Embedded Container Service bundled images.

4. Restart the Cloudera Embedded Container Service by selecting Cloudera Manager Clusters <<cluster name>> Actions Rolling Restart .

Figure 4: Rolling Restart



For OCP configuration settings

To enable a cluster-wide proxy, follow the instructions in the OpenShift Container Platform documentation: https://docs.openshift.com/container-platform/4.13/networking/enable-cluster-wide-proxy.html

Related Information

Updating proxy configuration in an existing workbench

Proxy setting best practices

Learn about best practices for using environment variables to propagate proxy settings.

In Kubernetes, proxy settings can be propagated to pods through the use of environment variables in pod specific configurations or through ConfigMaps or Secrets that are mounted as volumes within pods. This allows the proxy settings to be passed down to individual containers within the pods. However, it is important to note that not all applications may automatically inherit these settings, and some may require additional configuration within the container image or application code to properly utilize the proxy settings.

It is important to note that the use of no_proxy or NO_PROXY environment variables to bypass proxy settings may not be consistently respected by all third-party libraries or applications. While some libraries or applications may automatically honor these settings, others may not.

Many popular libraries and frameworks in various programming languages, such as Python, Java and Node.js, have their own way of handling proxy settings, which may not necessarily rely on the no_proxy or NO_PROXY environment variables. These libraries may have their own configuration files or internal settings that dictate how they handle proxy configurations, and these settings may not always align with the no_proxy or NO_PROXY environment variables set at the system or project level.

As a result, it is important to be aware that relying solely on no_proxy or NO_PROXY environment variables may not provide consistent results across all libraries or applications used in a project. In some cases, it may be necessary to clear out the http_proxy, https_proxy, no_proxy, or NO_PROXY environment variables in your project's environment variables or configuration files to ensure that the third-party libraries or applications do not attempt to apply proxy settings at all.

To ensure that proxy settings are consistently respected across all libraries and applications used in a project, it is recommended to carefully review the documentation and configuration options of each library or application, and configure them accordingly.

However, if you ever have to manage a stack written in multiple languages, you might need to consider some best practices for setting HTTP proxy configurations:

For http_proxy and https_proxy:

- Use lowercase. HTTP_PROXY is not always supported or recommended.
- If you absolutely must use the uppercase as well, be sure both versions share the same value.

For no_proxy:

- Use lowercase.
- Use comma-separated hostname:port values.
- IP addresses are acceptable, but hostnames are never resolved.
- Suffixes are always matched (for example, example.com will match test.example.com).
- If top-level domains need to be matched, avoid using a leading dot (.).
- Avoid using CIDR matching since only the Go and Ruby languages support that.

PIP installs with Proxy

When using pip to install packages from external sources through a proxy, it is generally recommended to perform this installation in a separate session (project session), utilizing the existing proxy rules. Proxy settings, including proxy server addresses, usernames, passwords, and other configurations, are typically environment-specific. By using a separate session, you can ensure that the appropriate proxy settings are used for the specific installation task. This helps to avoid conflicts or misconfigurations with your main session's proxy settings, which may be required for other tasks or applications.

To resolve Python communication issues between pods with proxy setup, in Project Settings Advanced add the four environmental vatiables (HTTP_PROXY, http_proxy, HTTPS_PROXY, https_proxy) set to empty values. This allows Python to run and use kubedns and kubeproxy properly.

Environment variables in the customer environment:

- HTTP_PROXY
- http_proxy
- HTTPS_PROXY
- https_proxy
- NO_PROXY
- no_proxy
- ALL_PROXY

The workaround of adding these environment variables in the project session helps resolve the proxy issue. Since Python does not support CIDR blocks for no_proxy, the request automatically gets directed to http_proxy or htpps_proxy, which causes the failure and prevents jobs from completing.

Related Information

We need to talk: Can we standardize NO_PROXY?
Installing a non-transparent proxy in a Cloudera AI environment

NTP proxy setup on Cloudera Data Engineering

Cloudera Data Engineering requires specific proxy configurations to manage virtual cluster connections efficiently in an air-gapped setup with restricted outbound connections. This setup ensures seamless access to external resources while adhering to network security and management policies.

Depending on your cluster platform, whether it is RKE2 for Cloudera Embedded Container Service clusters or OpenShift Container Platform, specific configurations and deployment methods apply. Consider the detailed instructions for configuring proxy settings in Cloudera Embedded Container Service server and agent configurations, as well as enabling cluster-wide proxies in OpenShift Container Platform environments. Also consider the provided sample configurations and guidelines for configuring proxy servers, including specifying Classless Inter-domain Routing (CIDR) ranges to exclude from proxy routing and updating proxy server allowlists.

Installing a non-transparent proxy in a Cloudera Data Engineering environment

If Cloudera Data Engineering is used in an air-gapped environment, a proxy configuration is not mandatory. If a non-transparent proxy is used, then certain endpoints must be added to the list of allowed endpoints for the proxy.

Before you begin

Configure the No Proxy value with the Classless Inter-Domain Routing (CIDR) ranges for the Nodes, POD CIDR, and Service CIDR. Any IP range for internal services with seamless internal network connectivity must be added in the No Proxy configuration. Specify these CIDR ranges in the configuration to ensure that the traffic destined for these ranges bypasses the proxy. Add comma-separated no-proxy configurations without any spaces between them.

About this task

If your Cloudera on premises deployment uses a non-transparent network proxy, configure proxy hosts that the workloads can use for connections with Cloudera Data Engineering virtual clusters. You can configure the proxy configuration values from the Cloudera Management Console.

The procedure for updating these settings might be different and dependent on the proxy server software used.



Important:

Apply these procedures before installing or upgrading the Cloudera Data Engineering virtual cluster.



Note: The configured settings apply for Cloudera on premises deployment using the Cloudera Embedded Container Service. In an OpenShift Container Platform deployment, the default values are used.

Procedure

1. In the Cloudera console, click the Cloudera Management Console tile. The Cloudera Management Console Home page displays.
2. On the left navigation menu, click Administration Networks tab. The **Networks** page displays.
3. Configure the following options for the proxy values:

Table 3: Proxy values

Field	Description
HTTPS proxy	It is the HTTP or HTTPS proxy connection string used with the Cloudera Data Engineering virtual clusters. You must specify this connection string in the form: http(s)://[***USERNAME***]:[***PASSWORD***]@[***HOST***]:[***PORT***]. The [***USERNAME***] and [***PASSWORD***] parameters are optional. You can specify the connection proxy string without these parameters.
HTTP proxy	It is the HTTP or HTTPS proxy connection string used with the Cloudera Data Engineering virtual clusters. You must specify this connection string in the form: http(s)://[***USERNAME***]:[***PASSWORD***]@[***HOST***]:[***PORT***]. The [***USERNAME***] and [***PASSWORD***] parameters are optional. You can specify the connection proxy string without these parameters.
No proxy	This is a comma-separated list of hostnames, IP addresses, or hostnames and IP addresses that should not be accessed through the specified HTTPS or HTTP proxy URLs. In case of Cloudera Embedded Container Service deployments, you must include no-proxy URLs for the following: • Hostnames and IP addresses of all the Cloudera Embedded Container Service base clusters in your deployment. • All Cloudera Base on premises cluster nodes. • CIDR IP addresses for internal operations in the ECS cluster: 10.42.0.0/16 and 10.43.0.0/16 • localhost, .local, .svc, and .svc.cluster.local • Kubernetes service and Pod CIDR blocks. For these CIDR ranges in Cloudera Embedded Container Service, open the Cloudera Embedded Container Service configuration page in Cloudera Manager for the Cloudera Embedded Container Service cluster and search for “cidr”. In case of OpenShift Container Platform deployments, you must include no-proxy URLs for the following: • Hostnames and IP addresses of all the OpenShift Container Platform base clusters in your deployment. • All Cloudera Base on premises cluster nodes. • IP addresses and Hostnames of the master node(s) of the Kubernetes cluster. • localhost, .local, .svc, and .svc.cluster.local • Kubernetes service and Pod CIDR blocks. For Pod CIDR ranges in OpenShift Container Platform, see CIDR range definitions. **Important:** In an air-gapped environment, Cloudera Data Engineering requires console.domain.entry.com in the no_proxy settings to work. Without this entry, client configurations do not download, and Cloudera Data Engineering Service or Virtual Cluster creations fail.

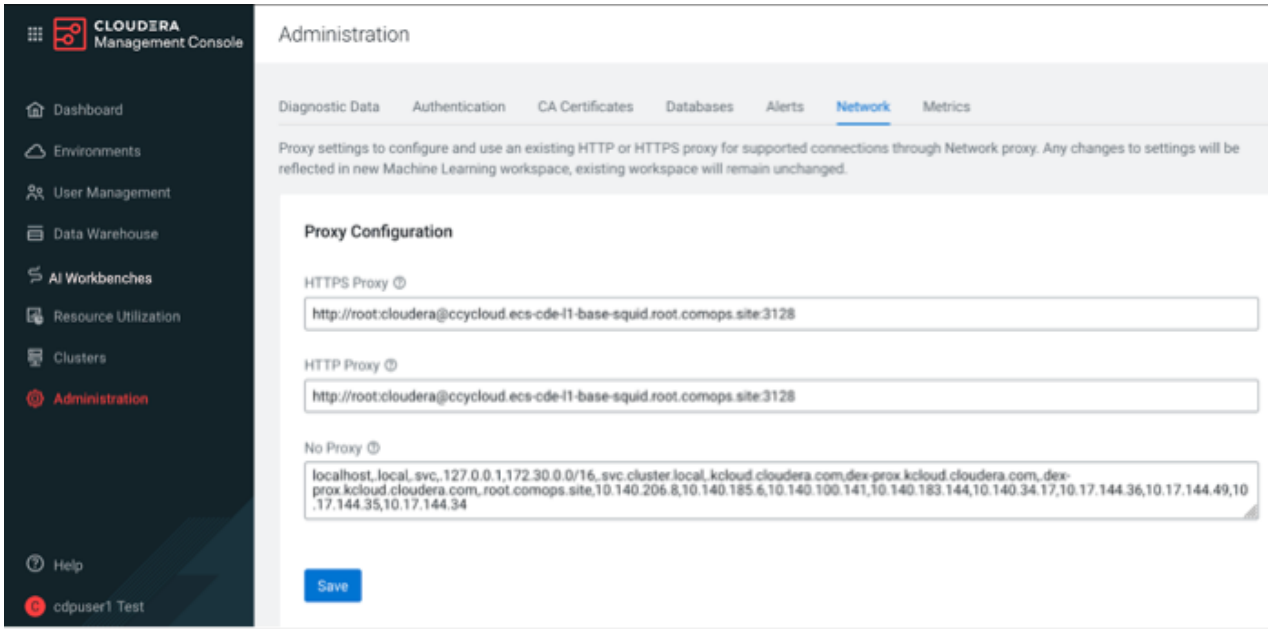
4. Click Save.

Important: In Cloudera Data Services on premises 1.5.5 SP3 and higher releases, if you update proxy configurations after creating Cloudera Data Engineering Service or Virtual Cluster, you must update and synchronize these changes in the Cloudera Data Engineering Service. For more information, see Updating and synchronizing stale configurations.

Example

Consider the following example:

Figure 5: NTP Proxy configuration example



Enabling proxies in Cloudera Embedded Container Service and OpenShift Container Platform environment

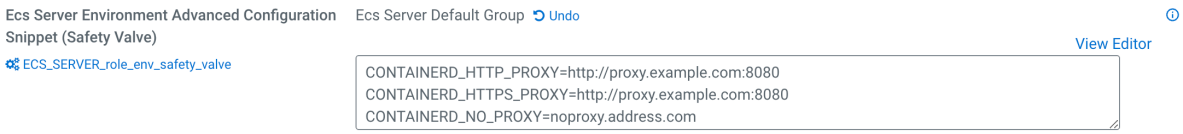
Depending on your cluster platform, whether it is RKE2 for Cloudera Embedded Container Service clusters or OpenShift Container Platform, specific configurations and deployment methods apply.

For ECS advanced configuration settings

1. Open Cloudera Manager Clusters Configuration Configuration Search .
2. Enter the name of the configuration you are looking for: ECS Server Environment Advanced Configuration.

   With the help of the Settings menu, include the environment variables in uppercase in the ECS Server Environment Advanced Configuration snippet (Safety Valve) for config.yaml to configure the CONTAINERD_HTTP_PROXY, CONTAINERD_HTTPS_PROXY and CONTAINERD_NO_PROXY.

   Figure 6: Environment variable settings in ECS Server Environment Advanced Configuration

   Example:

   CONTAINERD_HTTP_PROXY=http://proxy.example.com:8080

   CONTAINERD_HTTPS_PROXY=http://proxy.example.com:8080

   CONTAINERD_NO_PROXY=noproxy.address.com

   Note: Avoid using colons, double quotes, or white spaces as delimiters.

Note:

Ensure that CONTAINERD_NO_PROXY contains all the same hosts as the Cloudera on premisesconfiguration.

Include the host that holds the docker registry that serves all the normal Cloudera Embedded Container Service bundled images.

3. Enter the name of the configuration you are looking for: ECS Agent Environment Advanced Configuration.

   With the help of the Settings menu, include uppercase environment variables in the ECS Agent Environmen t Advanced Configuration snippet (Safety Valve) for config.yaml to configure the CONTAINERD_HTTP_ PROXY, CONTAINERD_HTTPS_PROXY and CONTAINERD_NO_PROXY.

   Figure 7: Environment variable settings in ECS Agent Environment Advanced Configuration

   Example:

   CONTAINERD_HTTP_PROXY=http://proxy.example.com:8080

   CONTAINERD_HTTPS_PROXY=http://proxy.example.com:8080

   CONTAINERD_NO_PROXY=noproxy.address.com

   Note: Avoid using colons, double quotes, or white spaces as delimiters.

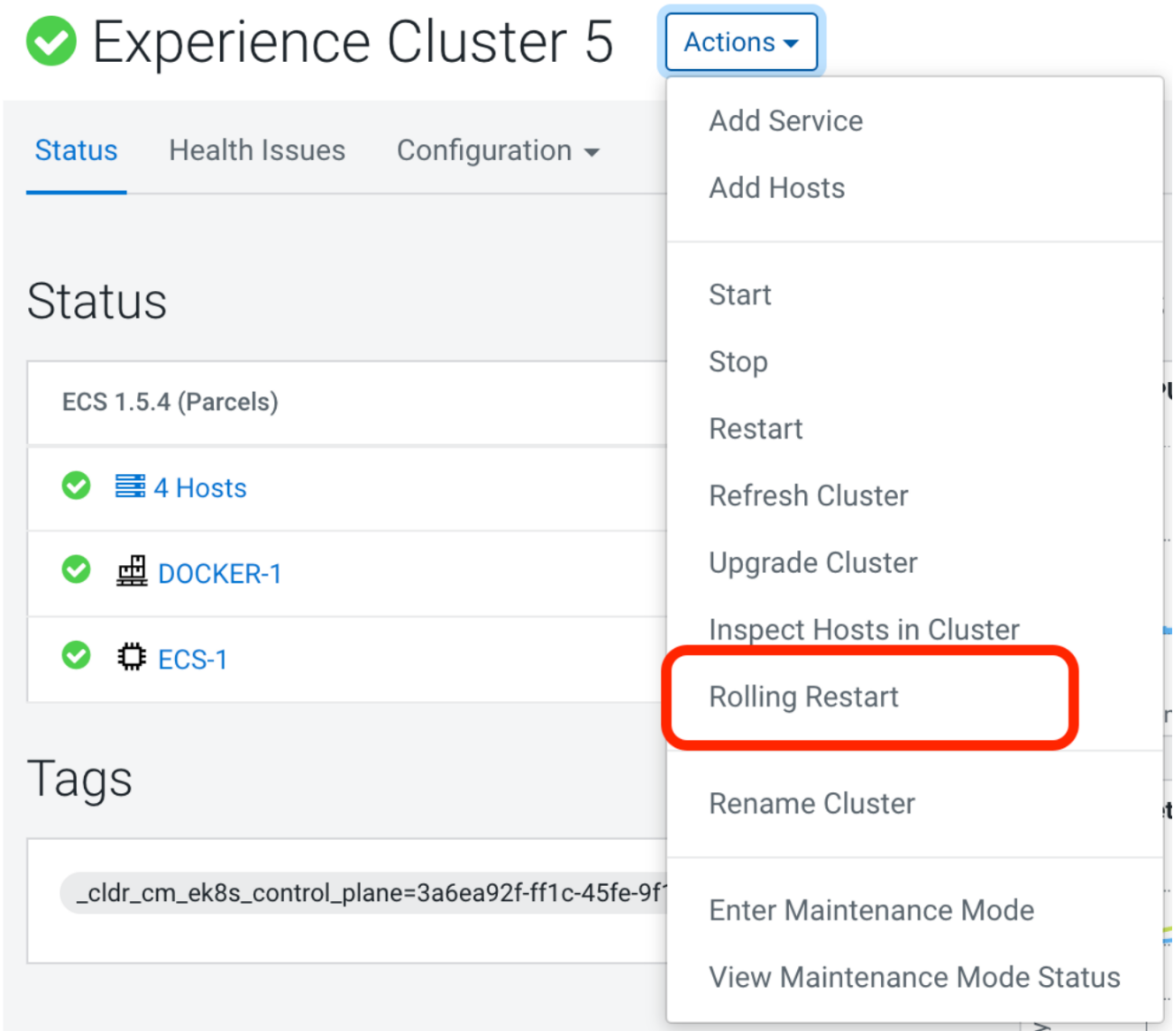
   Note:

   Ensure that CONTAINERD_NO_PROXY contains all the same hosts as the Cloudera on premises configuration.

   Pay special attention to include the host that holds the docker registry that serves all the normal Cloudera Embedded Container Service bundled images.

4. Restart the Cloudera Embedded Container Service by selecting Cloudera Manager Clusters *[***CLUSTER NAME***]* Actions Rolling Restart .

   Figure 8: Rolling Restart



For OCP configuration settings

To enable a cluster-wide proxy, follow the instructions in the OpenShift Container Platform documentation: https://docs.openshift.com/container-platform/4.13/networking/enable-cluster-wide-proxy.html

Related Information

Updating proxy configuration in an existing Cloudera Data Engineering

Proxy setting best practices

Learn about best practices for using environment variables to propagate proxy settings.

In Kubernetes, proxy settings can be propagated to pods through the use of environment variables in pod specific configurations or through ConfigMaps or Secrets that are mounted as volumes within pods. This allows the proxy settings to be passed down to individual containers within the pods. However, it is important to note that not all applications may automatically inherit these settings, and some may require additional configuration within the container image or application code to properly utilize the proxy settings.

It is important to note that the use of no_proxy or NO_PROXY environment variables to bypass proxy settings may not be consistently respected by all third-party libraries or application codes. While some libraries or application codes may automatically honor these settings, others may not.

Many popular libraries and frameworks in various programming languages, such as Python, Java and Node.js, have their own way of handling proxy settings, which may not necessarily rely on the no_proxy or NO_PROXY environment variables. These libraries may have their own configuration files or internal settings that dictate how they handle proxy configurations, and these settings may not always align with the no_proxy or NO_PROXY environment variables set in the Cloudera Management Console.

As a result, it is important to be aware that relying solely on no_proxy or NO_PROXY environment variables may not provide consistent results across all libraries or application codes used in a project. In some cases, it may be necessary to clear out the http_proxy, https_proxy, no_proxy, or NO_PROXY environment variables in your project's environment variables or configuration files to ensure that the third-party libraries or application codes do not attempt to apply proxy settings at all.

To ensure that proxy settings are consistently respected across all libraries and application codes used in a job or session, you must review the documentation and configuration options of each library or application code, and configure them accordingly.

However, if you ever have to manage a stack written in multiple languages, you might need to consider some best practices for setting HTTP proxy configurations:

For http_proxy and https_proxy:

- Use lowercase form. HTTP_PROXY is not always supported or recommended.
- If you absolutely must use the uppercase form as well, be sure both versions share the same value.

For no_proxy:

- Use lowercase form.
- Use comma-separated hostname:port values.
- IP addresses are acceptable, but hostnames are never resolved.
- Suffixes are always matched (for example, example.com will match test.example.com).
- If top-level domains need to be matched, avoid using a leading dot (.).
- Avoid using CIDR matching since only the Go and Ruby languages support that.

Related Information

We need to talk: Can we standardize NO_PROXY?

Installing a non-transparent proxy in a Cloudera Data Engineering environment