

Cloudera Streams Messaging Operator for Kubernetes 1.7.0

Cloudera Surveyor for Apache Kafka Operations

Date published: 2024-06-11

Date modified: 2026-07-31

CLOUDERA

<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

Get started with the Cloudera Surveyor for Apache Kafka UI.....	4
Entity summaries.....	8
Tables and filtering.....	10
Entity Details drawer.....	12
Alerts.....	13
Data refresh on the UI.....	15
Clusters.....	17
Downloading client configurations.....	18
Brokers.....	20
Editing broker properties.....	21
Topics.....	23
Creating topics.....	24
Editing topic properties.....	25
Deleting topics.....	27
Consumer Groups.....	28
Editing consumer group offsets.....	29
Deleting consumer groups.....	31
Producers.....	31

Get started with the Cloudera Surveyor for Apache Kafka UI

The Cloudera Surveyor UI provides a comprehensive set of features to manage and monitor Kafka entities such as clusters, brokers, topics, and consumer groups across dedicated overview and details pages.

The UI organizes functionality into dedicated pages accessible from the left-hand navigation, where you can manage core Kafka entities including clusters, brokers, topics, and consumer groups. When an instance name is configured for the deployment, it is displayed in the header to identify the Cloudera Surveyor instance.

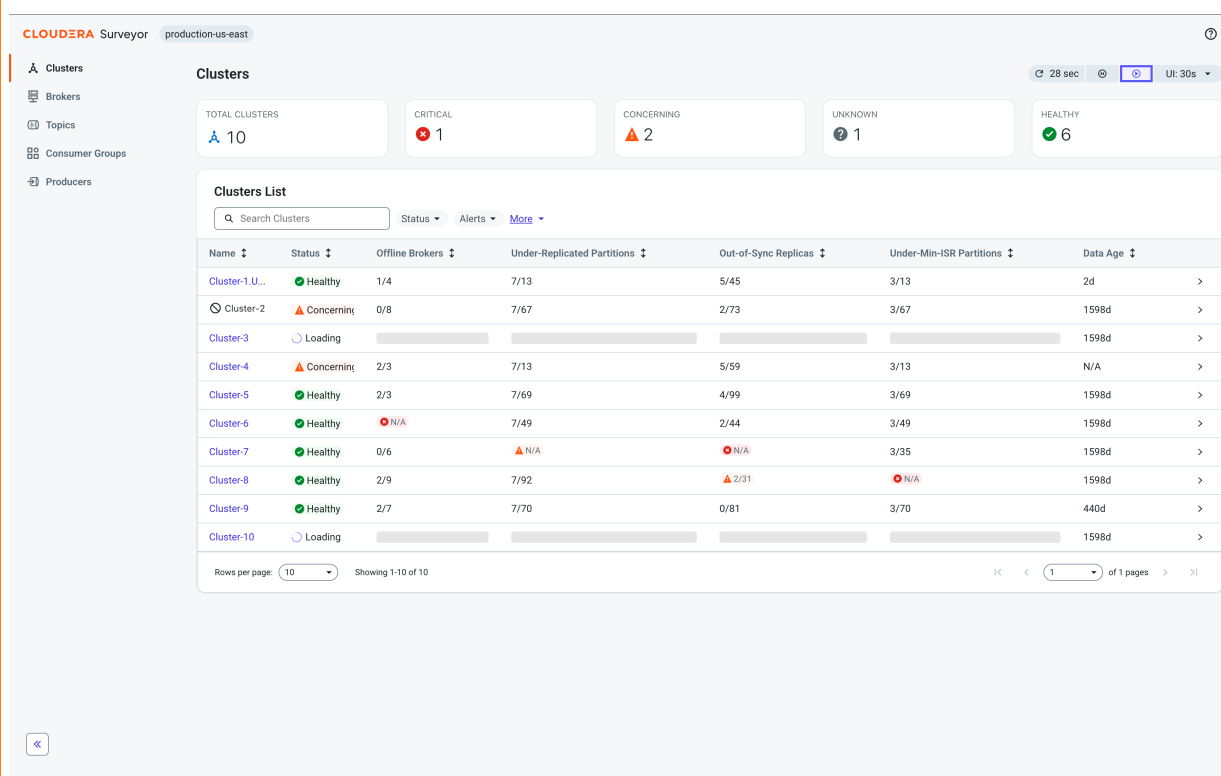
Each left-hand navigation item leads to an overview page, which displays an overview of all entities of that type. From an overview page, clicking an entity name opens the corresponding details page, which provides in-depth information about the selected individual entity.

Alert and status icons such as , , , and  are used throughout the UI to surface entity health, and highlight issues and errors at a glance.

For Clusters

View and monitor all configured Kafka clusters with overall health status and alert summaries. Filter and sort clusters by status, name, or alert severity. Click any cluster to open the **Cluster Details** page for comprehensive information including broker status, topic metrics, and cluster-wide configurations.

Figure 1: Clusters page




For Brokers

View overall status, health, disk usage, alert information, and more for all brokers within a selected cluster. Filter and sort brokers by status, ID, disk usage, or alert severity. Click any broker to open the **Broker Details** page for

in-depth information about broker status, hosted replicas, log directories, and configuration settings. Configure broker properties using the  Edit Properties option.


Figure 2: Brokers page



For Topics

View and manage all topics within a selected cluster, including status, partition count, replica distribution, and consumer group associations. Filter and sort topics by name, status, partition count, or alert severity. Click any topic to open the **Topic Details** page for comprehensive information about partitions, replicas, consumer groups, and topic configurations. Create new topics with custom configurations, update the configuration of existing topics, or delete existing topics.

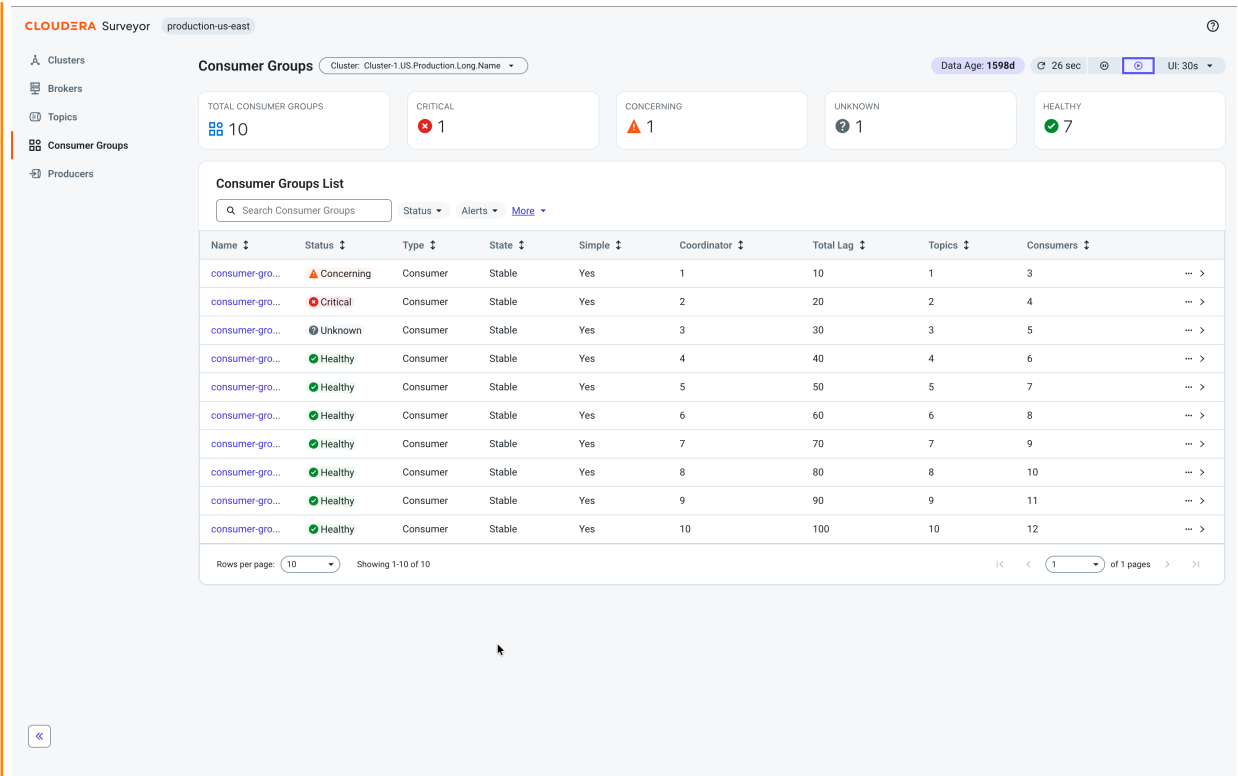
Figure 3: Topics page



For Consumer Groups

View and monitor all consumer groups within a selected cluster, including group status, member count, and lag metrics. Filter and sort consumer groups by name, status, group type, or lag. Click any consumer group to open the **Consumer Group Details** page for detailed information about partition assignments, members, subscribed topics, and offset positions. Delete consumer groups as needed using the  Delete option.


Figure 4: Consumer Groups page



For Producers

View and monitor all producers within a selected cluster, including producer status, partition count, and transaction details. Filter and sort producers by status, ID, or alert severity. Click any producer to open the **Producer Details** page for comprehensive information about alerts, transactions, and partition producer state.

Figure 5: Producers page

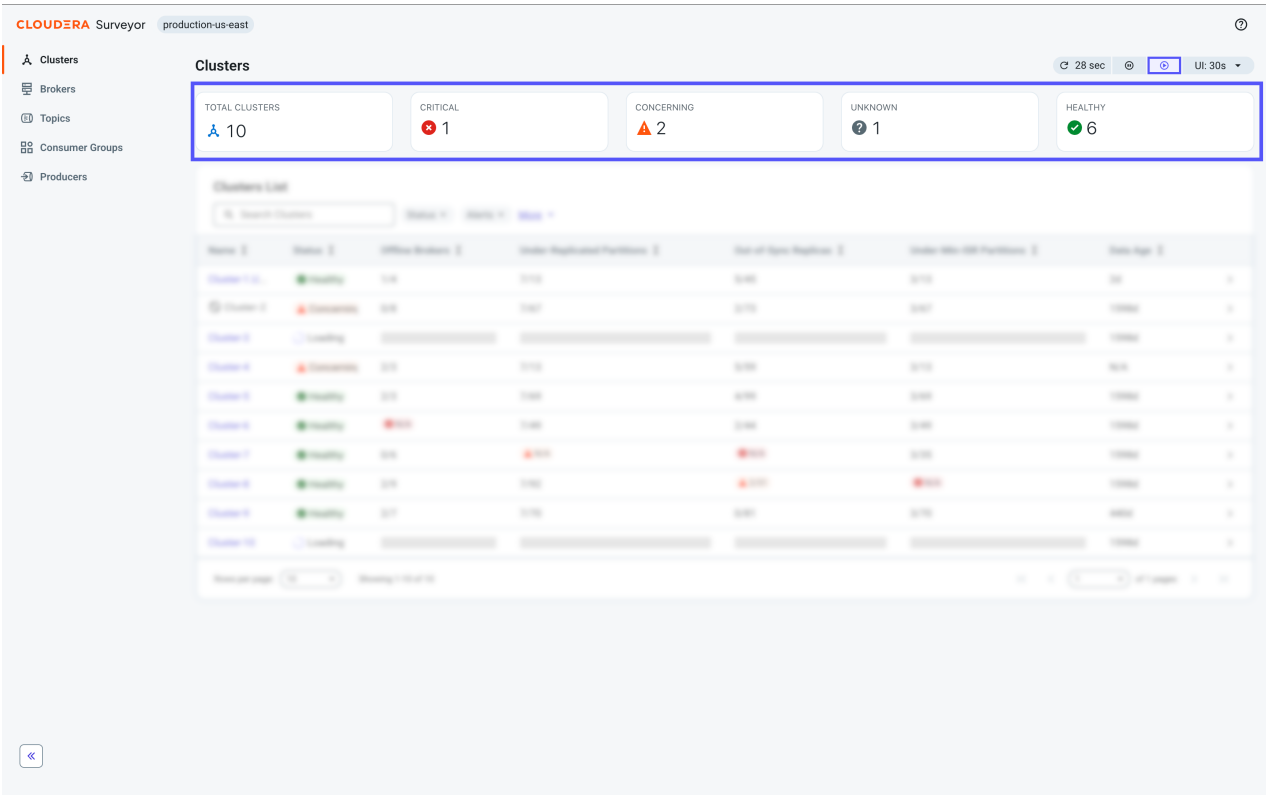


Entity summaries

All pages display a summary section at the top providing an at-a-glance view of entity health and important metrics without needing to scroll through detailed tables.

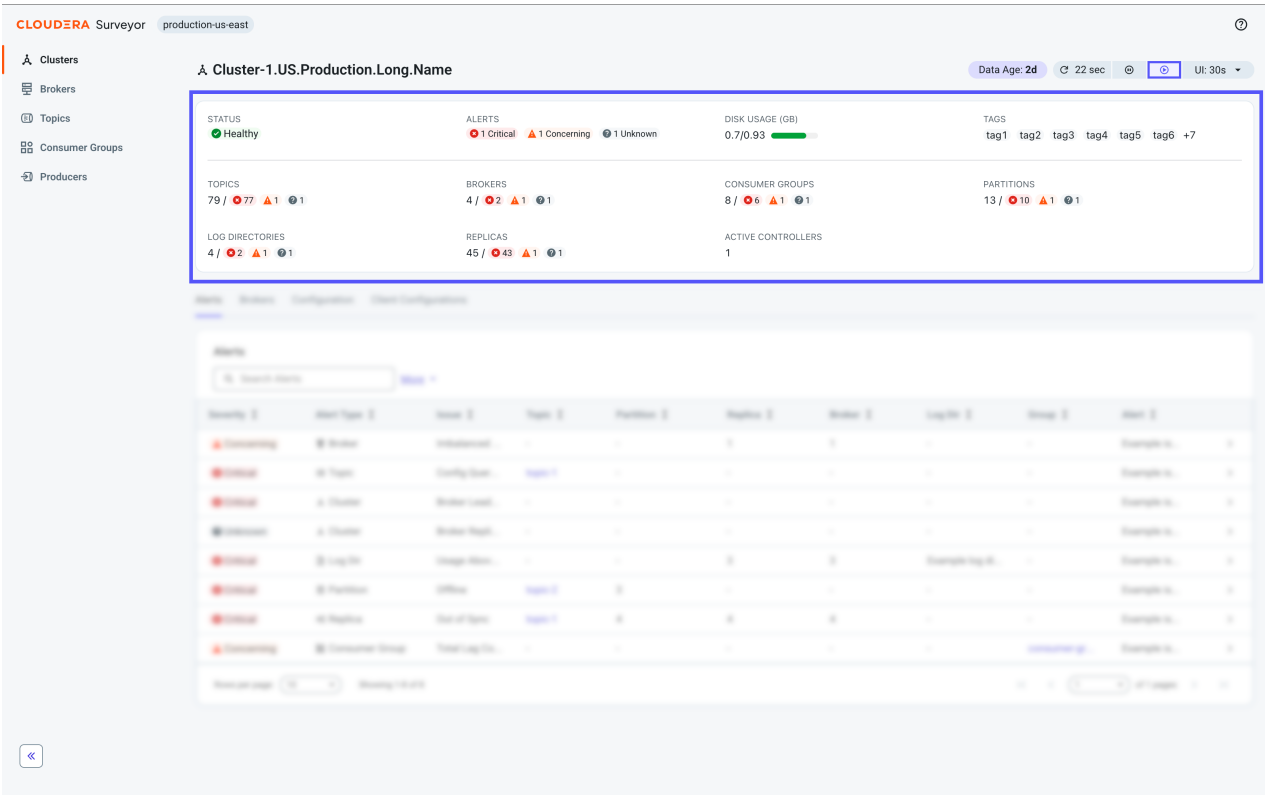
On overview pages the summary displays the total count of entities and a health breakdown displaying how many entities are in each status category. Each status category is displayed in a separate tile with an icon and count.

Figure 6: Clusters summary



On entity details pages, the summary displays key attributes specific to that entity in two sections. The upper section typically includes the entity's overall status, a count of active alerts by severity, and entity-specific attributes such as disk usage for clusters, state for consumer groups, or topic size for topics. Beneath a dividing line, a lower section provides aggregate counts and summaries of related entities. For example, the **Cluster Details** page displays total counts of related entities such as topics, brokers, consumer groups, and partitions.

Figure 7: Cluster details summary



Tables and filtering


Tables in Cloudera Surveyor support combinable filters displayed as pills in the filter bar, removable individually or all at once. Active filters persist in the URL, enabling you to bookmark and share specific filtered views.

Filters and controls

All tables include filters and filtering controls above the table. Overview pages typically provide the same set of quick filter drop-down lists for common filtering needs, while details pages and their tabs offer page-specific options depending on the data being displayed.

For example, the **Clusters** page includes the following filters:

Figure 8: Clusters page filters




- **Search box** – Free-text search that filters by entity name.
- **Status drop-down list** – Quick filter for entity status. Multiple statuses can be selected at once.
- **Alerts drop-down list** – Quick filter to display only entities with specific alert severities. Only one severity can be selected at a time.
- **More drop-down list** – Custom filtering with full control over options and filter operations.

Custom filtering

Access custom filtering using the More drop-down list. Select an option from the list, then choose a filter operation and specify your criteria. The available options are tailored to each page or tab and generally include a mix of visible table columns and additional entity attributes. Some visible table columns might not be available as filter options.

Figure 9: Example with multiple filters



Filter operations

Filter operations depend on the data type of the option you selected. The following operations are available per data types:

- **Text fields** (such as names, IDs, hosts):
 - Equals / Not Equal To – Exact match
 - Contains / Does Not Contain – Partial text match

- Matches Regular Expression – Pattern matching using regex
- In / Not In – Match any value from a list
- **Numeric fields** (such as counts, sizes, offsets):
 - Equals / Not Equal To – Exact match
 - Greater Than / Greater Than or Equal To – Compare numeric values
 - Less Than / Less Than or Equal To – Compare numeric values
 - In / Not In – Match any value from a list
- **Enumerated fields** (such as status, severity, type):
 - Equals / Not Equal To – Exact match
 - In / Not In – Match any value from a list
- **All fields:**
 - Is Null – Match null or missing values

Entity Details drawer

The Entity Details drawer provides quick access to all Kafka entity attributes in a searchable format organized into sections, without navigating away from the overview table.

Tables for entities such as brokers, topics, and consumer groups display the most important attributes for quick scanning, but each entity has many additional attributes. The **Entity Details** drawer provides access to all available attributes for the selected entity. Open the drawer by clicking anywhere in a row except the entity name to view the details of that entity.

Figure 10: Entity Details drawer for brokers

The screenshot shows the Cloudera Surveyor interface for a Kafka cluster. The left sidebar contains navigation links for Clusters, Brokers, Topics, Consumer Groups, and Producers. The main area displays the 'Brokers' section for a specific cluster. At the top, there are summary cards for 'TOTAL BROKERS' (3), 'CRITICAL' (0), 'CONCERNING' (1), 'UNKNOWN' (1), and 'HEALTHY' (1). Below this is a 'Brokers List' table with columns: ID, Short Name, Status, Alerts, Controller, Leaders, Replicas, and Disk Usage (GB). The first row is highlighted, corresponding to 'broker-1.kafka.local.test.short-name'. To the right of the table, the 'Entity Details' drawer is open, showing a search bar and several expandable sections: 'About' (Status: Healthy), 'Health' (Alerts: 1, Disk Usage: 50/100, Leader Replicas Log Size: 51200.00 MB), and 'Network' (Hostname: broker-1.kafka.local, Port: 9092, Address: broker-1.kafka.local:9092, Rack: -).

The drawer organizes attributes into named sections that vary by entity type. For example, the **Entity Details** drawer for brokers has About, Health, Network, and Associated Entities sections. You can expand or collapse each section

independently. Use the search field at the top of the drawer to filter attributes by name across all sections. Click any other row to switch the drawer to that entity.

Click the  icon to collapse the drawer to a narrow strip. After you collapse the drawer, the strip remains visible on the right edge of the page until you navigate away. The strip displays the name of the selected entity. While the drawer is collapsed, clicking a row updates the selected entity shown in the strip but does not expand the drawer.

Click the  icon at the bottom of the strip to expand the drawer again.

Alerts

Alerts surface problems detected in your Kafka environment, allowing you to quickly identify and investigate issues across monitored Kafka entities.

Cloudera Surveyor continuously evaluates the state of each monitored Kafka entity during every snapshot. When a condition outside expected bounds is detected, an alert is generated for that entity. Each alert has the following properties:

- **Severity** – Indicates the priority of the alert, categorized as  Critical,  Concerning, or  Unknown.
- **Alert type** – Identifies the entity type the alert belongs to, such as  Cluster,  Broker,  Topic,  Partition,  Log Dir.
- **Issue type** – Describes what was detected.
- **Cause** – Provides additional details.

Some alerts are always active by default and require no manual configuration, such as a broker going offline or a partition losing its leader. Other alerts are threshold-based, such as consumer group lag or broker disk usage, and only trigger when thresholds are configured. You can set these thresholds on a per-cluster basis in Cloudera Surveyor configuration. A comprehensive list of active alerts for your environment is available in the **Alerts** tab on each details page.

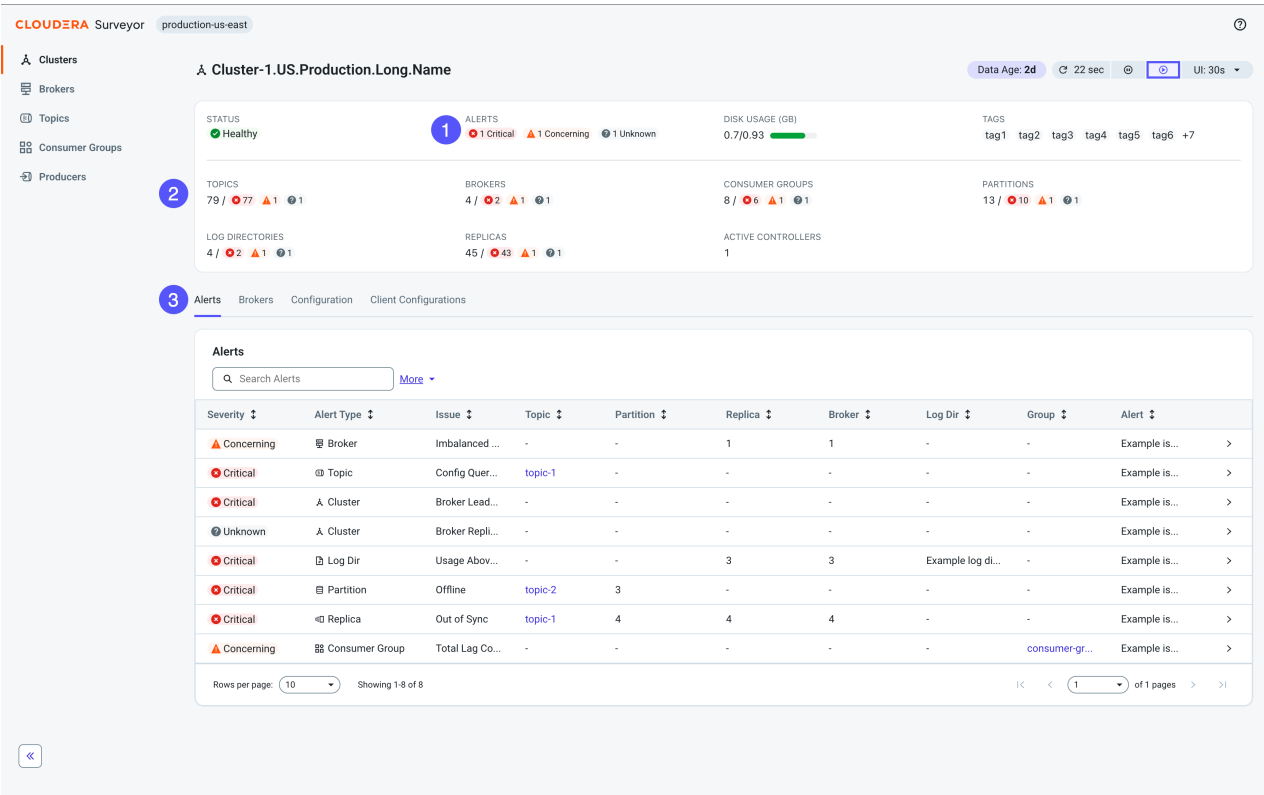
Alerts across the UI

Alerts can surface in the following locations throughout the UI:

- On overview pages – The **Status** column reflects each entity's health based on its alerts, and the **Alerts** quick filter lets you filter the table to display only entities with alerts of a specific severity.
- On details pages – An alert summary is displayed at the top of the page with the count of active alerts by severity. A dedicated **Alerts** tab provides the full list of alerts for that entity, including alerts from its child entities.
- On individual fields in tables and the **Entity Details** drawer – Alert indicators appear inline next to specific data values displaying when a particular metric has an issue.

The following example shows alerts on the **Cluster Details** page:

Figure 11: Alert example on Cluster Details page



- Alert summary displays the total count of alerts by severity for this cluster and all its child entities.
 - Child entity summaries display the total count and status distribution of child entities.
-  **Note:** The summaries display entity status counts, not alert counts. However, because each entity's status is determined by its active alerts, these indicators reflect the overall health impact of alerts across related entities.
- The **Alerts** tab displays the full list of all alerts with detailed information including severity, alert type, issue type, affected entities, and cause messages.

Alerts on individual fields

Alerts appear as inline indicators next to individual data fields throughout the UI, making problems visible at a glance wherever the affected data is displayed.

When Cloudera Surveyor detects a problem, it surfaces inline indicators on the affected fields. These indicators surface alerts that are also tracked in the **Alerts** tab. This visibility ensures you can quickly spot problems while scanning tables or viewing entity details, without navigating to the **Alerts** tab. These indicators are displayed as icons alongside field values in tables, detail pages, and the **Entity Details** drawer.

Figure 12: Fields with alerts

Clusters List

Search Clusters

Status Alerts More

Name	Status	Offline Brokers	Under-Replicated Partitions	Out-of-Sync Replicas	Under-Min-ISR Partitions	Data Age
Cluster-1.U...	Healthy	1/4	7/13	5/45	3/13	2d
Cluster-2	Concerning	0/8	7/67	2/73	3/67	1598d
Cluster-3	Loading					1599d
Cluster-4	Concerning	2/3	7/13	5/59	3/13	N/A
Cluster-5	Healthy	2/3	7/69	4/99	3/69	1599d
Cluster-6	Healthy	N/A	7/49	Offline	3/49	1599d
Cluster-7	Healthy	0/6	N/A	N/A	3/35	1599d
Cluster-8	Healthy	2/9	7/92	2/31	N/A	1599d
Cluster-9	Healthy	2/7	7/70	0/81	3/70	440d
Cluster-10	Loading					1599d

Rows per page: 10

Showing 1-10 of 10

1 of 1 pages

The following examples represent common field-level issues detected within your Kafka environment:

- Replication factors below the recommended minimum
- In-synch replica counts that do not match expectations
- Consumer group lag exceeding thresholds
- Partition sync states indicating offline replicas


Data refresh on the UI

Data refresh controls throughout the Cloudera Surveyor UI enable you to manage data update frequency, and provide information on how fresh the data is.

Data refresh controls

The Cloudera Surveyor UI includes various data refresh control options. These provide information on how recently the data was published in the UI, how old the latest snapshot is, and allow you to control UI update frequency. These controls do not offer a way to configure the snapshot interval or offer a way to trigger a new snapshot on-demand.

Figure 13: Data refresh controls



The following controls are available on all Cloudera Surveyor UI pages:

- Data Age: [TIME] – Snapshot time counter. Displays the time since Cloudera Surveyor took a snapshot of the Kafka cluster. Snapshots are taken automatically at preconfigured intervals.
-  [TIME] – Refresh timer and button. Displays the time until next automatic refresh to display latest snapshot data. Click to refresh the UI and display latest available snapshot data on-demand. Does not trigger a new Kafka snapshot.
-  /  – Pause and resume buttons for automatic refresh.
- UI: [TIME] – Dropdown that sets the automatic refresh frequency for displaying snapshot data.



Note: The Data Age: [TIME] counter is not displayed on the **Clusters** page, as snapshot age is only relevant within a specific cluster.

Data flow example

This examples shows a scenario in which snapshots are taken every 10 minutes:

1. Minute 0: Cloudera Surveyor takes a snapshot of the Kafka cluster
2. Minute 0-10:
 - Data Age: [TIME] counter starts at 0 and counts up to 10 minutes.
 - The snapshot data is published on the UI either automatically with the next automatic refresh or when you click the  [TIME] button.
 - Changes in Kafka that happened after the snapshot was taken are not visible.
3. Minute 10:
 - Cloudera Surveyor takes the next snapshot.
 - Data Age: [TIME] resets to 0.

- Minute 10+: After the next UI refresh, data from the new snapshot is displayed, including the most recent Kafka changes.



Note: Broker and topic configuration data is an exception to the snapshot system. Configuration changes are retrieved on demand during page load rather than relying on snapshots. Consequently, neither the backend snapshot settings nor the UI refresh controls affect this data.

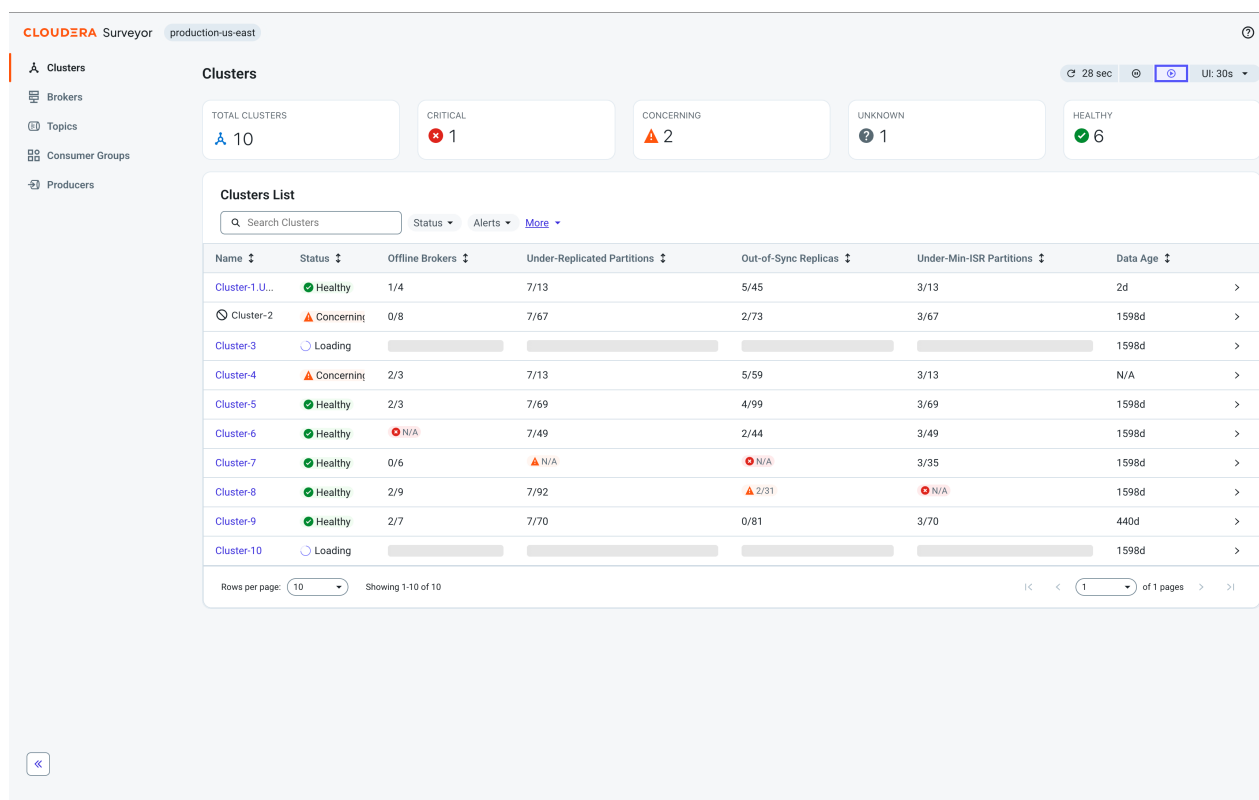
Clusters

Kafka cluster health, offline broker counts, under-replicated partitions, and out-of-sync replicas are monitored from the Clusters page in Cloudera Surveyor.

Clusters page

The **Clusters** page provides a comprehensive overview of all Kafka clusters monitored by Cloudera Surveyor. View cluster health status, active alerts, and key metrics at a glance. Filter and sort clusters by status, name, or alert severity to quickly identify clusters requiring attention.

Figure 14: Clusters page



At the top of the page, health summary tiles display the total number of clusters categorized by status, providing an at-a-glance view of the overall health distribution across all clusters.

Beneath the tiles, a table displays the following essential cluster metrics for quick health assessment:

- Name** – A unique identifier for the cluster. Click to open the **Cluster Details** page.
- Status** – The overall health status of the cluster.
- Offline Brokers** – The count of brokers that are offline.
- Under-Replicated Partitions** – The count of partitions with fewer replicas than configured.
- Out-of-Sync Replicas** – The count of replicas not in sync with their leaders.

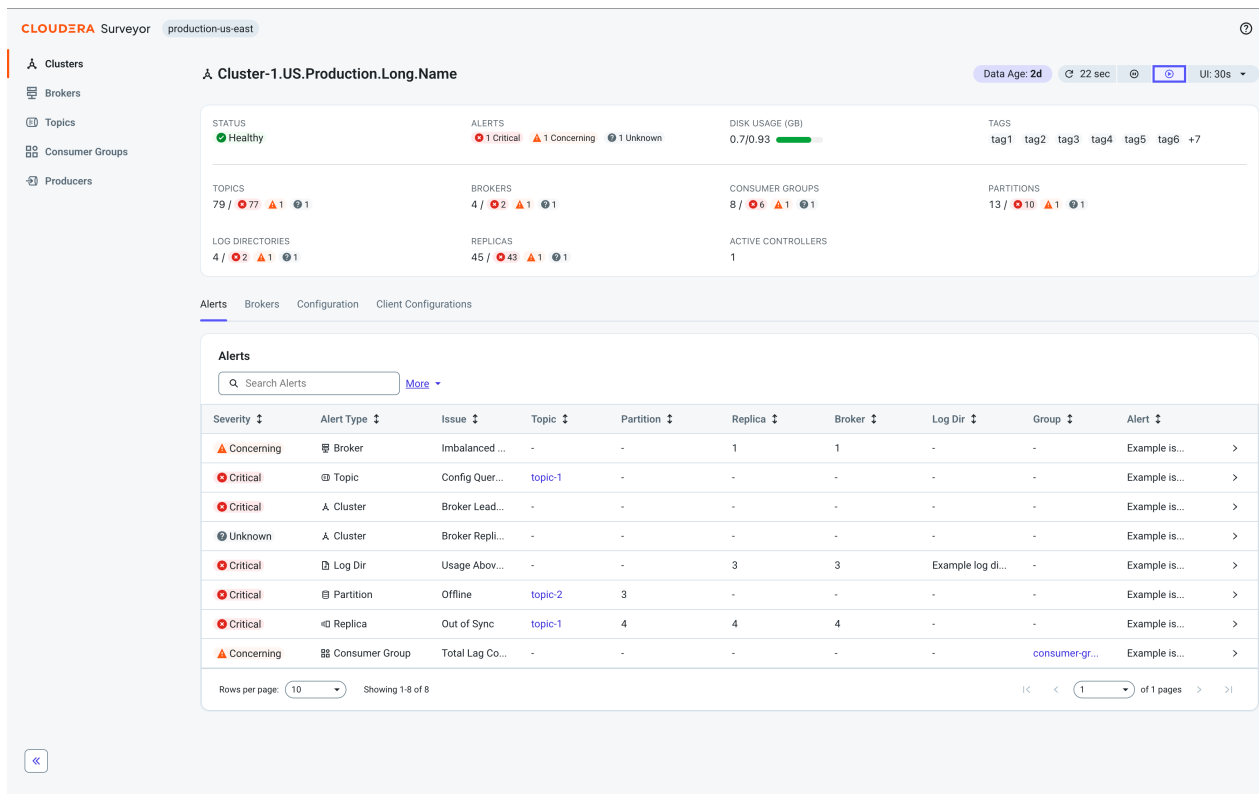
- **Under-Min-ISR Partitions** – The count of partitions with in-sync replicas below minimum.
- **Data Age** – The time elapsed since cluster data was last collected.

Click any cluster row to open the **Entity Details** drawer for quick access to all cluster attributes. Click the cluster name to go to the **Cluster Details** page.

Cluster Details page

The **Cluster Details** page provides comprehensive information about an individual cluster. At the top of the page, health summary tiles display the cluster's overall status and key metrics for child entities including brokers, topics, consumer groups, partitions, and replicas.

Figure 15: Cluster Details page



The page organizes cluster information through the following tabs:

- **Alerts** – View and monitor all alerts affecting the cluster.
- **Brokers** – View all brokers in the cluster and their status.
- **Configuration** – View aggregated configuration across all brokers in the cluster.
- **Client Configurations** – Preview, copy, and download client configuration files for the cluster, such as connection properties and CA certificates.

Downloading client configurations

Download client configuration files for a monitored Kafka cluster from Cloudera Surveyor to connect external clients to that cluster.

About this task

The **Client Configurations** tab contains a bundle of client configuration files for a Kafka cluster monitored by Cloudera Surveyor. When the tab is populated, it displays a bundle title, a file count, and a list of files with optional

descriptions. You can preview, copy, or download individual files. Use Download All to download the complete bundle as a zip archive.

The files available on this tab depend on how Cloudera Surveyor is configured. Typical bundles include a client configuration properties file, one or more CA certificates, and optionally a README or other instructional document. Your Cloudera Surveyor administrator configures what is available on this tab. For information about configuring the bundle contents, see [Exposing client configurations for download](#).

If the **Client Configurations** tab is empty, no files or download actions are available. This typically indicates that your Cloudera Surveyor administrator has not configured a client configuration bundle for this cluster. Contact your Cloudera Surveyor administrator.

The screenshot shows the Cloudera Surveyor interface for a cluster named "Cluster-1.US.Production.Long.Name". The left sidebar contains navigation links for Clusters, Brokers, Topics, Consumer Groups, and Producers. The main panel displays various cluster metrics: STATUS (Healthy), ALERTS (1 Critical, 1 Concerning, 1 Unknown), DISK USAGE (0.7/0.93 GB), TAGS (tag1 through tag6), TOPICS (79 / 77 / 1), BROKERS (4 / 2 / 1), CONSUMER GROUPS (8 / 6 / 1), PARTITIONS (13 / 10 / 1), LOG DIRECTORIES (4 / 2 / 1), REPLICAS (45 / 43 / 1), and ACTIVE CONTROLLERS (1). Below these metrics, the "Client Configurations" tab is active, showing a "Kafka Client Configuration" bundle with 4 files available for download. The files listed are: README.md (text/markdown - Usage instructions and required system property notes), external.properties (text/plain - Kafka client properties for the external listener), kafka-ca.crt (application/x-pem-file - CA certificate to trust the Kafka broker), and oauth-ca.crt (application/x-pem-file - CA certificate to trust the OAuth token endpoint). Each file has "Preview" and "Download" buttons. A "Download All" button is also present at the top right of the file list.

Procedure

1. In Cloudera Surveyor, go to Clusters and select your Kafka cluster.
2. Open **Cluster Details** and select the Client Configurations tab.
3. Download files from the bundle.
 - To download all files as a zip archive, click Download All.
 - To download an individual file, locate it in the list and click Download next to it.
4. To review a file before downloading, click Preview or Copy next to it.

What to do next

If you downloaded the full bundle, extract the archive to a working directory. Keep all files together in the same directory because some configuration files might reference other files in the bundle by relative path.

Related Information

[Exposing client configurations for download](#)


Brokers

Kafka broker status, disk usage, leader and replica counts, and configuration properties are monitored and edited from the Brokers page in Cloudera Surveyor.

Brokers page

The **Brokers** page provides detailed information about all brokers within a selected cluster. Monitor broker health, resource usage, and active alerts. Filter and sort brokers by status, ID, disk usage, or alert severity to quickly identify brokers requiring attention.

Figure 16: Brokers page



At the top of the page, health summary tiles display the total number of brokers categorized by status, providing an at-a-glance view of the overall health distribution across all brokers in the selected cluster.

Beneath the tiles, a table displays the following essential broker metrics for quick health assessment:

- **ID** – A unique identifier for the broker. Click to open the **Broker Details** page.
- **Short Name** – A shortened broker hostname for easier identification. Calculated by removing the longest common suffix from broker hostnames, or using hostname:port when hostnames are identical.
- **Status** – The overall health status of the broker.
- **Alerts** – A summary of alert counts by severity.
- **Controller** – An indication of whether this broker is the cluster controller. This column is only displayed for ZooKeeper-based clusters.
- **Leaders** – The number of partition leaders hosted on the broker.
- **Replicas** – The number of partition replicas hosted on the broker.
- **Disk Usage (GB)** – The total disk space used across all log directories.

Click any broker row to open the **Entity Details** drawer for quick access to all broker attributes. Click the broker ID to go to the **Broker Details** page.

Broker Details page

The **Broker Details** page provides in-depth information about an individual broker's status and health. At the top of the page, health summary tiles display the broker's overall status and key metrics including disk usage, replica count, and leader count.

Figure 17: Broker Details page

The screenshot displays the Cloudera Surveyor interface for a specific broker. The top navigation bar includes 'CLUSTERA Surveyor' and 'production-us-east'. The left sidebar shows a hierarchy: Clusters, Brokers (selected), Topics, Consumer Groups, and Producers. The main content area is titled 'Cluster: Cluster-1.US Production.Long.Name' and shows a 'Data Age: 937d' and '28 sec' refresh rate. The broker's status is 'Healthy'. Key metrics are displayed in a grid: TOPICS (24), REPLICAS (192), LEADERS (64), and LEADER REPLICAS LOG SIZE (51200.00 MB). Below this, there are tabs for Alerts, Configuration, Log Directories, and Replicas. The Alerts tab is active, showing a search bar and a table of alerts. The table has columns for Severity, Alert Type, Issue, Broker ID, Short Name, and Alert. Two alerts are listed: a 'Critical' alert for 'Log Dir' with the issue 'Usage Above Concern...' and a 'Concerning' alert for 'Broker' with the issue 'Offline'. The bottom of the page shows pagination: 'Rows per page: 10' and 'Showing 1-2 of 2'.

The page organizes broker information through the following tabs:

- **Alerts** – View and monitor all alerts affecting the broker.
- **Configuration** – View and edit broker configuration properties.
- **Log Directories** – Monitor log directory status and disk usage.
- **Replicas** – View all partition replicas hosted on this broker.

Editing broker properties

Edit broker properties in Cloudera Surveyor using the  Edit Properties option on the Configuration tab of the Broker Details pages. This option provides comprehensive control over broker settings.



Important: If an external tool manages your Kafka cluster, that system might control configuration. Changes made directly through Cloudera Surveyor might be overwritten during the subsequent configuration reconciliation or deployment cycle.

About this task

Broker configuration properties control Kafka broker behavior and settings. In Cloudera Surveyor, you modify these properties using the  **Brokers** **Broker Details Configuration**  **Edit properties** option, which opens the **Edit Properties** modal. The modal displays all broker configuration properties retrieved from Kafka and allows you to bulk edit with a single operation. Modified properties are marked with indicators, and changes are staged for review before being applied. If you navigate away with unsaved changes, warnings are displayed and your edits are automatically preserved.

The property list includes both static and dynamic Kafka properties. However, you can only successfully modify dynamic properties.

Whether a property is dynamic or static is indicated by the **Read Only** attribute. Static properties are marked as **Read Only: yes**, while dynamic properties are **Read Only: no**. When you attempt to modify a static property and apply changes, Kafka rejects the modification. This is because static properties are fundamental broker settings that require a broker restart to take effect. Because Cloudera Surveyor does not manage the lifecycle of clusters and cannot restart brokers, you cannot change these properties through the UI.

Before you begin

- You must have the necessary permissions to modify broker configurations in your Kafka cluster. The required Kafka ACL is `ALTER_CONFIGS` on the `BROKER` resource.
- Broker configuration editing must be enabled for the target cluster. If the **Edit Properties** button is greyed out, editing is disabled for this cluster through Cloudera Surveyor configuration. For more information, see [Disabling broker configuration editing](#).

Procedure

1. In Cloudera Surveyor, click  **Brokers**.
 2. Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
 3. Click the ID of a broker to open the **Broker Details** page for that broker.
 4. Go to the **Configuration** tab.
 5. Click  **Edit Properties**.
 6. Edit properties in the **Edit Properties** modal.
 - a) Find the properties you want to modify.
 - b) Select or deselect the checkbox next to each property to add or remove property overrides.
 - c) Update the property values using the Property Value field.
 - d) Click OK.
-  **Tip:** Selecting a checkbox adds the property override even if the value matches the default.
7. Review your changes and click  **Apply Changes**.
Properties that are modified are marked with an indicator pill. You can also choose to discard your changes by clicking  **Discard Changes**.
 8. In the confirmation dialog, click **Apply** to confirm your changes.

Results

Your configuration changes are immediately applied to Kafka and are instantly visible in Cloudera Surveyor. While the configuration values update immediately in Cloudera Surveyor, propagation of these changes to all brokers might take a few moments.

Topics

Kafka topic health, partition distribution, replica placement, and consumer group subscriptions are monitored and managed from the Topics page in Cloudera Surveyor.

Topics Page

The **Topics** page provides comprehensive information about all topics in the selected cluster. Monitor topic health, partition distribution, replica status, and consumer group subscriptions. Filter and sort topics by name, status, partition count, or alert severity. Create new topics or delete existing ones as needed.

At the top of the page, health summary tiles display the total number of topics categorized by status, providing an at-a-glance view of the overall health distribution across all topics in the selected cluster.

Beneath the tiles, a table displays the following essential topic metrics for quick health assessment:

- **Name** – A unique identifier for the topic. Click to open the **Topic Details** page.
- **Status** – The overall health status of the topic.
- **Under-Replicated** – The count of partitions with fewer in-sync replicas than configured.
- **Offline Partitions** – The count of partitions with no active leader.
- **Retention Period** – The duration for which messages are retained.
- **Replication Factor** – The number of replicas for each partition.
- **Topic Size (GB)** – The total size of all replicas.



Tip:

The **Topics** page includes a Visibility quick filter to filter between visible and hidden topics. This filter is unique to the **Topics** page. Topics are marked as hidden based on a configurable name pattern.

The default pattern is: `^(__,+|ATLAS_,+)`. This pattern filters topics that use common internal topic naming conventions. You can customize this pattern in Cloudera Surveyor configuration.

Click Add Topic to create a new topic with custom configurations. Click any topic row to open the **Entity Details** drawer for quick access to all topic attributes. Click the topic name to go to the **Topic Details** page.

Topic Details Page

The **Topic Details** page provides comprehensive information about an individual topic. At the top of the page, health summary tiles display the topic's overall status and key metrics including partition count, replication factor, and partition health indicators.

The screenshot shows the Cloudera Surveyor interface for a topic named 'topic-1'. The left sidebar contains navigation links for Clusters, Brokers, Topics (selected), Consumer Groups, and Producers. The main content area displays various metrics for the topic:

- STATUS:** Healthy (green checkmark)
- ALERTS:** 1 Unknown (yellow warning icon)
- TOPIC SIZE:** N/A
- CLEANUP:** N/A
- RETENTION PERIOD:** N/A
- PARTITIONS:** 10
- REPLICAS:** 30
- MIN ISR:** 2
- REPLICATION FACTOR:** 3
- COMPRESSION:** gzip
- CONSUMER GROUP LAG:** 1 / 3 / 5

Below the metrics, there are tabs for Alerts, Configuration, Partitions, Replicas, Consumer Groups, and Producers. The 'Alerts' tab is active, showing a search bar and a table of alerts:

Severity	Alert Type	Issue	Partition	Replica	Broker	Alert
Critical	Partition	Under Replicated	1	-	-	Several partitions ...
Concerning	Replica	Out of Sync	5	1	1	In-sync replica cou...

At the bottom of the alerts table, it shows 'Rows per page: 10' and 'Showing 1-2 of 2'. There are also pagination controls at the bottom right.

The page organizes topic information through the following tabs:

- **Alerts** – View and monitor all alerts affecting the topic.
- **Configuration** – View and edit topic configuration properties. Increase partition count and manage property overrides.
- **Partitions** – Monitor partition health and replica distribution.
- **Replicas** – View detailed information regarding the partitions of this topic.
- **Consumer Groups** – View consumer groups consuming from this topic and monitor their lag.
- **Producers** – View producers associated with this topic.

Creating topics

Create new topics with custom properties in Cloudera Surveyor using the  Add Topic option on the Topics page.

About this task

Topics are the fundamental units of data organization in Kafka, organizing messages into named streams that producers write to and consumers read from. In Cloudera Surveyor, you create topics using the  Topics  Add Topic option, which opens the  Create Topic drawer. The drawer provides an interface to configure mandatory

properties as well as property overrides. Property overrides customize behaviors such as data retention or cleanup policies, overriding the default values defined at the server level.

Choose your partition count and replication factor carefully. Partition count can be increased later, but cannot be decreased. Changing the replication factor is possible, but requires partition reassignment, which is not possible using Cloudera Surveyor. For more information, see [Increasing replication factor](#) in the Apache Kafka documentation.

Before you begin

- You must have the necessary permissions to create topics in your Kafka cluster. The required Kafka ACL is CREATE on the CLUSTER or TOPIC resource.

Procedure

- In Cloudera Surveyor, click  Topics.
- Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
- Click  Add Topic at the top right of the table.
- In the  Create Topic drawer, enter the following required information:
 - Topic Name** – A unique identifier for the topic.
 - Partitions** – The number of partitions (determines parallelism).
 - Replication Factor** – The number of replicas for each partition. Replicas cannot exceed the broker count.
- Configure property overrides in the **Property Overrides** section.
 - Click  Edit Properties.
 - In the **Edit Properties** modal, find the properties you want to modify.
 - Select the checkbox next to each property to include it in your changes.
 - Update the property values using the Property Value field.
 - Click OK.

 **Tip:** You can also edit individual properties that are added to the **Property Overrides** section by clicking the  icon or remove them by clicking the  icon. You can also clear all overrides using the  Clear All option.
- Review your configuration and click  Create Topic.
- In the confirmation dialog, click Create to confirm topic creation.

Results

The topic is created in Kafka. However, the new topic will not be visible in Cloudera Surveyor until the next Kafka snapshot is taken and data from the new snapshot is published to the UI. The  Create Topic drawer remains open, the Name field is cleared and focused. All other fields and property overrides are preserved, so you can immediately create another topic without reopening the drawer.

Editing topic properties

Edit topic configuration properties in Cloudera Surveyor using the  Edit Properties option on the Configuration tab of the Topic Details pages.



Important: If an external tool manages your Kafka cluster, that system might control configuration. Changes made directly through Cloudera Surveyor might be overwritten during the subsequent configuration reconciliation or deployment cycle.

About this task

Topic configuration properties control various aspects of topic behavior including data retention, cleanup policies, compression settings, and replication behavior. In Cloudera Surveyor, you modify these properties using the 

Topics Topic Details Configuration  Edit Properties option, which opens the Edit Properties modal. The modal displays all topic configuration properties retrieved from Kafka and allows you to bulk edit with a single operation. Modified properties are marked with indicators, and changes are staged for review before being applied. If you navigate away with unsaved changes, warnings are displayed and your edits are automatically preserved. The Configuration tab also provides an option to increase the partition count.

When you modify topic configurations, the changes take effect immediately without requiring a broker restart. However, the replication factor of a topic is read-only and cannot be modified through the Cloudera Surveyor UI. Replication factor is set when the topic is created. To change the replication factor of an existing topic, you must use Kafka's partition reassignment tool. For more information, see [Increasing replication factor](#) in the Apache Kafka documentation.

Before you begin

- You must have the necessary permissions to modify topic configurations in your Kafka cluster. The required Kafka ACL is ALTER_CONFIGS on the TOPIC resource.

Procedure

- In Cloudera Surveyor, click  Topics.
- Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
- Click the Name of a topic to open the **Topic Details** page for that topic.
- Go to the **Configuration** tab.
- Increase partition count.
 - Locate the **PARTITIONS** field at the top of the tab.
 - Click the  icon next to the field.
 - Enter a new value that is greater than the existing value.



Important: You cannot decrease the partition count.

- Click Save.
- Edit topic properties.
 - Click  Edit Properties.
 - In the **Edit Properties** modal, find the properties you want to modify.
 - Select or deselect the checkbox next to each property to add or remove property overrides.
 - Update the property values using the Property Value field.
 - Click OK.



Tip: Selecting a checkbox adds the property override even if the value matches the default.

- Review your changes and click  Apply Changes.

Properties that are modified are marked with an indicator pill. You can also choose to discard your changes by clicking  Discard Changes.

8. In the confirmation dialog, click Apply to confirm your changes.

Results

Your configuration changes are immediately applied to Kafka and are instantly visible in the Cloudera Surveyor UI.

Deleting topics

Delete topics in Cloudera Surveyor using the  Delete option on the Topics pages or the  Delete Topic option on the Topic Details pages.



Important: If an external tool manages your Kafka cluster, topics deleted directly through Cloudera Surveyor might be recreated during the subsequent reconciliation or deployment cycle.

About this task

Deleting a topic permanently removes all messages, partition data, and topic metadata from the cluster. This operation cannot be undone.

You delete topics when decommissioning services, cleaning up test or development topics, or removing obsolete data streams. Deleting a topic immediately disrupts any applications that are still using the topic.

Before you begin

- You must have the necessary permissions to delete topics in your Kafka cluster. The required Kafka ACL is DELETE on the TOPIC resource.
- You must verify that all consumers and producers have stopped using the topic.
- You must verify that no applications, system, data pipeline, or other tooling depends on the topic.

Procedure

1. In Cloudera Surveyor, click  Topics.
2. Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
3. Locate the topic you want to delete in the table.
4. Click   Delete on the right side of the topic row.



Tip: You can also delete a topic from the **Topic Details** page by clicking Actions  Delete .

5. Type delete in the confirmation dialog.
6. Click  Delete to confirm deletion.

Results

The topic and all its data are permanently deleted from Kafka. However, the deleted topic will remain visible in Cloudera Surveyor until the next Kafka snapshot is taken and data from the new snapshot is published to the UI.

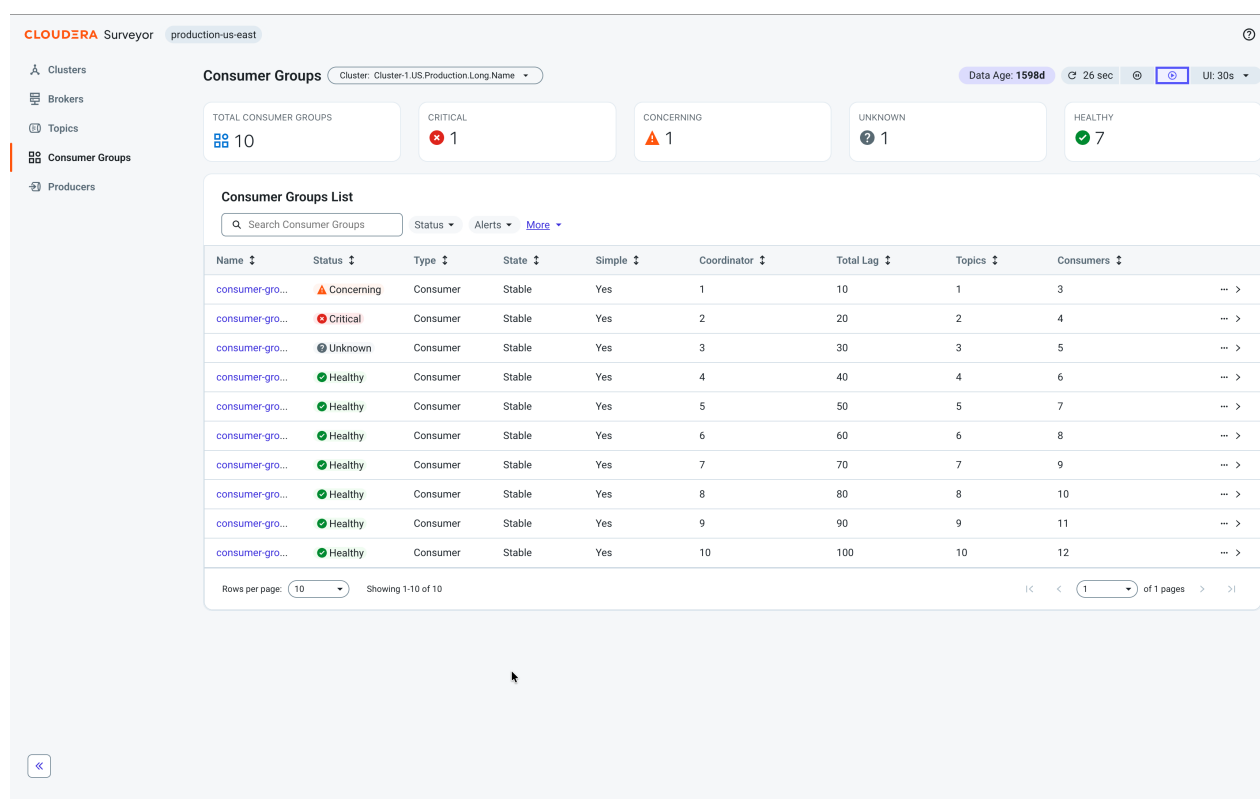
Consumer Groups

Kafka consumer group lag, partition assignments, member status, and committed offset positions are monitored and managed from the Consumer Groups page in Cloudera Surveyor.

Consumer Groups Page

The **Consumer Groups** page provides detailed information about all consumer groups in the selected cluster. Monitor consumer group health, lag metrics, member status, and partition assignments. Filter and sort consumer groups by name, status, group type, or lag. Delete consumer groups as needed.

Figure 18: Consumer Groups page



At the top of the page, health summary tiles displays the total number of consumer groups categorized by status, providing an at-a-glance view of the overall health distribution across all consumer groups in the selected cluster.

Beneath the tiles, a table displays the following essential consumer group metrics for quick health assessment:

- **Name** – A unique identifier (Group Id) for the consumer group.
- **Status** – The overall health status of the consumer group derived from active alerts.
- **Type** – The type of consumer group such as Consumer, Classic, or Streams.
- **State** – The current state of the consumer group such as Stable, Empty, or Dead.
- **Simple** – The indication of whether the consumer group is a simple consumer group. Simple consumer groups have committed offsets but do not use the entire consumer group protocol with coordinator or rebalancing.
- **Coordinator** – The broker ID of the group coordinator responsible for managing this consumer group.
- **Total Lag** – The combined lag across all partitions consumed by the consumer group.
- **Topics** – The number of topics with offsets committed by the consumer group.
- **Consumers** – The number of active consumer members in the group.

Click any consumer group row to open the **Entity Details Drawer** for quick access to all consumer group attributes. Click the consumer group name to go to the **Consumer Group Details** page.

Consumer Group Details Page

The **Consumer Group Details** page provides detailed information about an individual consumer group. At the top of the page, health summary tiles display the consumer group's overall status and key metrics including member count, total lag, and subscribed topic count.

Figure 19: Consumer Group Details page

The screenshot shows the Cloudera Surveyor interface for a consumer group named 'consumer-group-1'. The page is organized into several sections:

- Header:** Includes the Cloudera Surveyor logo, cluster name 'production-us-east', and a dropdown for the cluster. It also shows 'Data Age: 1599d', '23 sec', and a 'Delete' button.
- Summary Tiles:**
 - STATUS:** Concerning (indicated by a red triangle icon).
 - ALERTS:** 2 Concerning (red triangle icon), 1 Unknown (blue circle icon).
 - STATE:** STABLE.
 - TYPE:** CONSUMER.
 - CONSUMERS:** 3.
 - TOTAL LAG:** 10.
 - ASSIGNED TOPICS:** 1.
 - COORDINATOR:** 1.
- Tabs:** Alerts, Partition Assignment, Members. The 'Alerts' tab is currently selected.
- Alerts Section:**
 - A search bar labeled 'Search Alerts' and a 'More' link.
 - A table with columns: Severity, Alert Type, Issue, Consumer Group, and Alert.
 - Three alerts are listed:

Severity	Alert Type	Issue	Consumer Group	Alert
Concerning	Consumer Group	Total Lag Concerning	consumer-group-1	Example issue cause for ...
Critical	Consumer Group	Total Lag Critical	consumer-group-2	Example issue cause for ...
Critical	Consumer Group	Total Lag Critical	consumer-group-3	Example issue cause for ...
 - Footer: 'Rows per page: 10', 'Showing 1-3 of 3', and pagination controls.

The page organizes consumer group information through the following tabs:

- **Alerts** – View and monitor all alerts affecting the consumer group.
- **Partition Assignment** – View partition offset positions and lag. Edit offsets as needed.
- **Members** – View active consumer group members and their partition assignments.

Editing consumer group offsets

Edit consumer group offsets in Cloudera Surveyor using the Edit Committed Offsets option on the Partition Assignment tab of the Consumer Group Details pages.

About this task

Consumer offsets track the position of each consumer group within a topic partition, allowing consumers to resume reading from where they left off. In Cloudera Surveyor, you edit offsets using the  Consumer Groups Consumer Group Details Partition Assignment  Edit Committed Offsets option, which opens the Edit Committed Offsets modal. The modal allows you to select a scope and choose from multiple reset strategies to control where consumers will begin reading messages.

Editing offsets changes the position from which consumers read messages when they restart. This operation only updates the stored offset position in Kafka. Active consumers must be stopped to ensure that the updated offsets take effect.

You edit offsets to replay messages after a processing error, skip problematic messages that are causing consumer failures, or reprocess data with updated business logic. Offset edits are also useful for testing consumer behavior or recovering from incorrect offset commits.

Before you begin

- You must have the necessary permissions to alter consumer group offsets in your Kafka cluster. The required Kafka ACL is READ on the GROUP resource.
- You must stop all consumer applications in the target consumer group before editing offsets. Active consumers do not pick up the updated offsets until they restart and rejoin the consumer group.

Procedure

1. In Cloudera Surveyor, click  Consumer Groups.
2. Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
3. Click the Name of a consumer group to open the **Consumer Group Details** page for that consumer group.
4. Go to the **Partition Assignment** tab.
5. Click  Edit Committed Offsets to open the offset editor.
6. In the **Edit Committed Offsets** modal, select a scope from the Scope drop-down menu.
 - All Subscribed Topic Partitions – Apply the offset edit to all partitions consumed by this consumer group.
 - Topics – Select specific topics to edit offsets for all their partitions.
 - Topics & Partitions – Select specific topic partitions individually.
7. If you selected the Topics or Topics & Partitions scope, select the topics and partitions from the searchable list.
8. In the **Selected Strategy** section, select a strategy.

The following strategies are available:

 - Set To Current – Resets to the consumer group's currently committed offset.
 - Set To Earliest – Resets to the beginning of the partition (oldest available message).
 - Set To Latest – Resets to the current end of the partition (newest available message).
 - Set To Date Time – Resets to the offset at a specific date and time. When selected, a Date Time field is displayed to enter the date and time.
 - Set To Time Duration – Resets to the offset to a point in time based on a specified duration. When selected, a Duration (ms) field is displayed to enter the duration in milliseconds.
 - Set To Custom Offset – Resets to an exact offset value. When selected, an Offset field is displayed to enter the offset value.
9. Click Apply.

Results

The offset positions are updated in Kafka. When consumer applications restart, they resume reading from the updated positions.

What to do next

Restart consumer applications to begin reading from the updated offset positions.

Deleting consumer groups

Delete consumer groups in Cloudera Surveyor using the  Delete option on Consumer Groups pages or the  Delete option on the Consumer Group Details pages.

About this task

Deleting a consumer group permanently removes the group's metadata and all committed offset records from the cluster. This operation cannot be undone.

You delete consumer groups when decommissioning applications, cleaning up test or development environments, or removing stale consumer groups that are no longer in use. When a deleted consumer group is recreated by restarting an application, it starts fresh with no offset history, and the offset behavior is determined by the `auto.offset.reset` property (typically starting from latest or earliest).

Before you begin

- You must have the necessary permissions to delete consumer groups in your Kafka cluster. The required Kafka ACL is DELETE on the GROUP resource.
- You must verify that all consumer applications in the target group have stopped. Deleting an active consumer group might cause unexpected behavior when consumers attempt to commit offsets or rejoin.
- You must verify that no applications or monitoring tools depend on the consumer group's offset information.

Procedure

1. In Cloudera Surveyor, click  Consumer Groups.
2. Click the cluster name in the breadcrumb at the top left and select your target cluster from the list.
3. Locate the consumer group you want to delete in the table.
4. Click   Delete on the right side of the consumer group row.



Tip: You can also delete a consumer group from the **Consumer Group Details** page by clicking the  Delete option.

5. Type delete in the confirmation dialog.
6. Click  Delete to confirm deletion.

Results

The consumer group and its committed offsets are permanently deleted from Kafka. However, the deleted consumer group will remain visible in Cloudera Surveyor until the next Kafka snapshot is taken and data from the new snapshot is published to the UI.

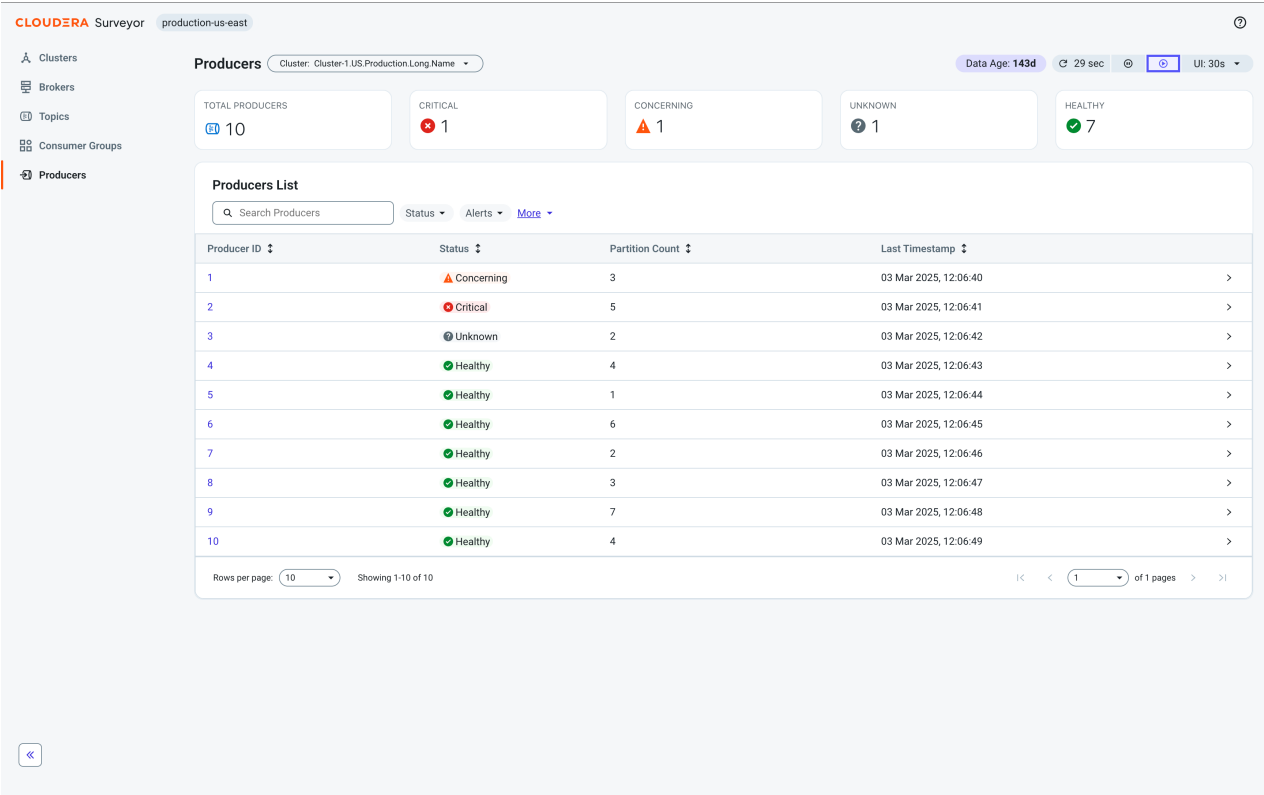
Producers

Kafka producer status, transaction state, and partition output are monitored from the Producers page in Cloudera Surveyor.

Producers page

The **Producers** page provides detailed information about all producers in the selected cluster. Monitor producer health, transaction status, and partition output. Filter and sort producers by status, ID, or alert severity to quickly identify producers requiring attention.

Figure 20: Producers page



At the top of the page, health summary tiles display the total number of producers categorized by status, providing an at-a-glance view of the overall health distribution across all producers in the selected cluster.

Beneath the tiles, a table displays the following essential producer metrics for quick health assessment:


- **Producer ID** – A unique identifier for the producer. Click to open the **Producer Details** page.
- **Status** – The overall health status of the producer.
- **Partition Count** – The number of partitions this producer is writing to.
- **Last Timestamp** – The most recent timestamp of producer activity.

Click any producer row to open the **Entity Details** drawer for quick access to all producer attributes. Click the producer ID to go to the **Producer Details** page.

Producer Details page

The **Producer Details** page provides comprehensive information about an individual producer's transactions and partition state. At the top of the page, health summary tiles display the producer's overall status, alert summary, and key metrics including partition count and last timestamp.

Figure 21: Producer Details page



The page organizes producer information through the following tabs:

- **Alerts** – View and monitor all alerts affecting the producer.
- **Transactions** – View all transactions associated with this producer and their state.
- **Partitions** – View partition producer state and epoch information for each partition.