



Confluent Cloud Service Level Agreement

This Confluent Cloud Service Level Agreement (“SLA”) describes the service availability commitment for the applicable Confluent Cloud Service purchased by Customer pursuant to the applicable agreement for the Cloud Service (“Agreement”) between the Confluent entity which entered into the Agreement (“Confluent”) and Customer. Unless otherwise provided herein, this SLA is subject to the terms of the Agreement and capitalized terms will have the meaning specified in the Agreement.

During Customer’s use of the Cloud Service, Confluent will use commercially reasonable efforts to make the Cloud Service available within the applicable Service Level. If Confluent does not meet the Service Level, Customer may be eligible to receive a Service Credit as described below. Confluent may update the SLA from time to time; provided, however, updates to the SLA will apply only to Orders entered into after the effective date of any update.

General Terms

Definitions

- **“Downtime”** is defined for the applicable Cloud Service function in the Service Specific Terms below. Downtime does not include unavailability that results from any of the exclusions set forth below or in the Service Specific Terms. Partial minutes of unavailability will not be counted as Downtime. Confluent’s monitoring system connects to the same endpoints that Customer uses.
- **“Monthly Uptime Percentage”** is defined for each applicable Cloud Service function in the Service Specific Terms below.
- **“Service Credit”** means the percentage of monthly fees attributable to Customer’s spend for the applicable Cloud Service function, calculated in Service Specific Terms, and credited to the Cloud Service bills in accordance with the process described in this SLA.
- **“Service Level”** means the Monthly Uptime Percentage for the applicable Cloud Service function as detailed in the Service Specific Terms section below.

Service Credits

Service Credits are calculated as a percentage of the total monthly fees paid by Customer for unavailable Cloud Service function for the calendar month in which the applicable Cloud Service function does not meet the Service Level, in accordance with the schedule below.

Service Credits are not refundable and can only be used toward future billing charges. Confluent will apply any Service Credits against Customer’s next billing charge. Service Credits are exclusive of any applicable taxes charged to Customer or collected by Confluent. Service Credits will not entitle Customer to any refund or other payment from Confluent. Service Credits are Customer’s sole and exclusive remedy for any unavailability of the Cloud Service in accordance with the terms of this SLA. Service Credits expire without refund twelve (12) months from issuance.



Service Credit Request and Application Process

To receive a Service Credit, Customer must submit a claim by logging a support ticket (if Customer is community supported, Customer must email cloud-support@confluent.io). To be eligible, the credit request must be received by Confluent within five (5) calendar days after the last day of the month in which the applicable Cloud Service function does not meet the Service Level, and must include all information reasonably necessary for Confluent to verify the claim, including:

1. the words "SLA Credit Request" in the subject line;
2. a description of the applicable client(s), the version of each such client, and the configurations for each such client; and
3. a description of the events resulting in Downtime, including the time and duration of the Downtime and Customer requests logs that document the failed write attempts.

Confluent will evaluate Customer requests and determine in good faith whether a Service Credit is owed based on its system logs, monitoring reports, configuration records, and other available information. If Confluent confirms that the Monthly Uptime Percentage applicable to the month of such request did not meet the Service Level, then Confluent will issue the Service Credit to Customer within one billing cycle following the month in which Customer's request is confirmed. Customer's failure to provide the request and other information as required above will disqualify Customer from receiving a Service Credit.

Exclusions

This SLA and any applicable Service Level does not apply to any unavailability or performance issues that results from:

1. A suspension described in Section 5.4 of the Agreement (Late Payments) or similar section; Customer's misuse of the Cloud Service in violation of Section 2 of the Agreement (Customer Use) or similar section; or Customer's violation of the Agreement or Confluent's Acceptable Use Policy;
2. Factors outside Confluent's reasonable control, including but not limited to any force majeure event, network intrusions, denial of service attacks, systemic internet issues or any other act or omission of any telecommunication or services provider;
3. Use of services, hardware, or software provided by a third party and not within the primary control of Confluent, including issues resulting from inadequate bandwidth or resulting from failures of cloud platform services on which the Cloud Service runs;
4. Customer's unauthorized action or lack of action when required, including those of Customer's Users or by means of Customer's passwords;
5. Customer's failure to use Confluent-supported Kafka clients with acceptable configuration values as defined in the Cloud Service documentation;
6. Customer-controlled actions and/or environment or other failures or shortcomings not within Confluent's control;
7. Failure by Customer to take any remedial action in relation to the Services as recommended by Confluent, or otherwise preventing Confluent from doing so;
8. Customer's negligence or willful misconduct, including failure to follow agreed-upon procedures;
9. Scheduled maintenance that takes place upon five (5) days email notice or ad hoc maintenance carried out to avoid future unavailability, and/or updates; or
10. Customer's failure to provide information required by Confluent to provision the Cloud Service.



Service Specific Terms

Confluent Cloud

Additional Definitions

- **“Downtime”** is the total accumulated minutes during a calendar month for a given Cloud Service cluster during which the entire cluster is unavailable. A minute is considered unavailable for a given cluster if all continuous attempts by Confluent’s monitoring system to write to the cluster within the minute fail.
- **“Monthly Uptime Percentage”** means the total number of minutes in a calendar month, minus the number of minutes of Downtime in such month, divided by the total number of minutes in such month. If Customer’s Cloud Service cluster is provisioned and running for only part of a calendar month, such cluster is deemed to be 100% available during the portion of the month in which it was not provisioned and running.

Service Credit				
Monthly Uptime Percentage	Basic Cluster	Standard (1 ECKU) or Enterprise (1 ECKU) Cluster	Dedicated (Single Zone) Cluster	Standard (2 ECKU), Enterprise (2 ECKU), or Dedicated (Multi-zone) Cluster
< 99.99% ≥ 99%	-	-	-	10%
< 99.95% ≥ 99%	-	-	10%	
< 99.9% ≥ 99%	-	10%	-	
< 99.5% ≥ 99%	10%	-	-	
< 99% ≥ 95%	25%	25%	25%	25%
< 95%	100%	100%	100%	100%

Confluent Cloud for Apache Flink®

Additional Definitions

- **“Confluent Cloud for Apache Flink”** means part of the Cloud Service that provides a cloud-native, serverless stream processing service based on Apache Flink, as further described in the Documentation.
- **“Downtime”** is the total accumulated one-minute periods during a calendar month during which Confluent Cloud for Apache Flink is unavailable in a given cloud region. A one-minute period is



considered unavailable for a given region if all continuous requests by Confluent’s monitoring system to either the applicable Flink Compute Pool API or Statement API result in a Fail Request.

- **“Fail Request”** means a request that returns an error code or otherwise does not return any successful code within one minute for the GET request due to an issue solely within Confluent’s control.
- **“Flink Compute Pool API”** means a software interface used to manage resources used to run Apache Flink workloads, as further described in the Documentation.
- **“Monthly Uptime Percentage”** means the total number of minutes in a calendar month, minus the number of minutes of Downtime in such month, divided by the total number of minutes in such month. If Customer’s Confluent Cloud for Apache Flink is provisioned and running for only part of a calendar month, such cluster is deemed to be 100% available during the portion of the month in which it was not provisioned and running.
- **“Statement API”** means a software interface used to model SQL statements for execution in Confluent Cloud for Apache Flink, as further described in the Documentation.

Monthly Uptime Percentage	Service Credit
< 99.99% ≥ 99.9%	5%
< 99.9% ≥ 99.0%	10%
< 99.0%	25%

Confluent Cloud Fully Managed Connectors

Additional Definitions

- **“Confluent Cloud Fully Managed Connectors”** are connectors that are managed and supported by Confluent in the Cloud Service and identified in the Documentation as “Supported”. Confluent Cloud Fully Managed Connectors exclude Connect with Confluent connectors, custom connectors, preview connectors, connectors uploaded by Customer to the Cloud Service, or connectors not suitable for production use, as set forth in the Documentation.
- **“Confluent Cloud Fully Managed Connector Instance”** is a deployment of one Confluent Cloud Fully Managed Connector type using one (1) or more Tasks.
- **“Downtime”** is the total accumulated consecutive five-minute periods during a calendar month for a given Confluent Cloud Fully Managed Connector Instance that is unavailable. The consecutive five-minute periods are calculated on a time-based tumbling window. A consecutive five-minute period is considered unavailable for a Confluent Cloud Fully Managed Connector Instance that is in a Failed State for the entire consecutive five-minute period. If a Confluent Cloud Fully Managed Connector Instance has failed due to a User Actionable Error, the Confluent Cloud Fully Managed Connector Instance is considered available.
- **“Failed State”** means the status of a Confluent Cloud Fully Managed Connector Instance where all Tasks for a Confluent Cloud Fully Managed Connector Instance have failed due to an issue solely within Confluent’s control.



- **“Monthly Uptime Percentage”** means the total number of consecutive five-minute periods in a calendar month, minus the number of consecutive five-minute periods of Downtime in such month, divided by the total number of consecutive five-minute periods in such month. If Customer’s Confluent Cloud Fully Managed Connector Instance is provisioned and running for only part of a calendar month, such instance is deemed to be 100% available during the portion of the month in which it was not provisioned and running.
- **“Task”** means a logical unit of work copying data.
- **“User Actionable Error”** represents a failure in any Task for the Confluent Cloud Fully Managed Connector Instance that is caused by the Customer. User Actionable Error includes an error that causes any Task for the Confluent Cloud Fully Managed Connector Instance to fail due to the Confluent Cloud Fully Managed Connector Instance being configured incorrectly by Customer, connections to external systems, throttling caused by Customer, and Customer server-side issues.

Service Credit			
Monthly Uptime Percentage	Confluent Cloud Fully Managed Connector Instance connecting to a Multi-zone Dedicated, Enterprise, or Standard Cluster	Confluent Cloud Fully Managed Connector Instance connecting to a Single-zone Dedicated or Standard Cluster	Confluent Cloud Fully Managed Connector Instance connecting to a Basic Cluster
< 99.99% ≥ 99.95%	10%	-	-
< 99.95% ≥ 99.5%		10%	10%
< 99.5% ≥ 99%			
< 99%	25%	25%	25%

Confluent Cloud KSQL

Additional Definitions

- **“Downtime”** is the total accumulated five-minute periods during a calendar month for a given ksqlDB Instance during which the entire ksqlDB Instance is unavailable. A five-minute period is considered unavailable for a given ksqlDB Instance if all Metadata Request attempts within the five-minute period fail.
- **“ksqlDB Instance”** is a multi-zone ksqlDB deployment using eight (8) or more Confluent Streaming Units.
- **“Metadata Request”** is any of the following requests as detailed in the Documentation: LIST STREAMS; LIST TABLES; or LIST QUERIES.
- **“Monthly Uptime Percentage”** means the total number of five-minute periods in a calendar month, minus the number of five-minute periods of Downtime in such month, divided by the total



number of five-minute periods in such month. If Customer’s ksqlDB Instance is provisioned and running for only part of a calendar month, such instance is deemed to be 100% available during the portion of the month in which it was not provisioned and running.

Monthly Uptime Percentage	Service Credit
< 99.9% ≥ 99.5%	5%
< 99.5% ≥ 99.0%	10%
< 99.0%	25%

Confluent Cloud Schema Registry

Additional Definitions

- **“Downtime”** is the total accumulated minutes during a calendar month for a given Confluent Cloud Service Schema Registry cluster during which the entire cluster is unavailable. A minute is considered unavailable for a given cluster if all continuous attempts by Confluent’s monitoring system to write to the cluster within the minute fail.
- **“Monthly Uptime Percentage”** means the total number of minutes in a calendar month, minus the number of minutes of Downtime in such month, divided by the total number of minutes in such month. If Customer’s Confluent Cloud Service Schema Registry cluster is provisioned and running for only part of a calendar month, such cluster is deemed to be 100% available during the portion of the month in which it was not provisioned and running.

Monthly Uptime Percentage	Stream Governance Essentials Service Credit	Stream Governance Advanced Service Credit
< 99.99% ≥ 99.5%	-	5%
< 99.5% ≥ 99.0%	5%	10%
< 99.0%	10%	25%