

Streaming Data Service by the Experts Behind Apache Kafka™

Confluent Cloud™ is a streaming data service that delivers the resilient, scalable streaming platform that is Apache Kafka. Deployable in minutes, it is the ideal streaming data service for the operations-starved organization or the agile, cloud-first developer on a mission. Confluent Cloud is a fully managed, private cluster service hosted on AWS (US-West-2, US-East-1), Google Cloud Platform, and Microsoft Azure*.

Built by the original creators of Apache Kafka from LinkedIn, Confluent Cloud operationalizes all of their knowledge and experience managing large-scale streaming data environments into a fully managed service makes implementing and managing an enterprise streaming platform with Kafka easy, reliable, and secure.

Why Confluent Cloud?



Unleash developer velocity

Simple, resilient, secure and performant, your team can focus on what matters most—velocity.



Set down the ops burden

The only thing harder than building and troubleshooting a distributed system is operating it at scale.



Use the Kafka ecosystem

Use Kafka APIs to leverage pre-built clients, connectors, and tools



You can change your mind

“Lift and shift” your Kafka applications from any location into—or out of—Confluent Cloud.

Not a Java Shop? No worries. We're polyglots.

When using one of the many Confluent clients, simply point the application at the Confluent Cloud and with a few simple configuration changes you'll be up and running. Python, C/C++, Go, and .NET all are supported.

Features:



Latest Kafka capabilities:

Always uses the latest stable Apache Kafka version. We take care of the upgrades so you don't have to.



Performant:

Highest throughput rate of any streaming data service. Standard plan is 5 MB/s. Custom plans available for any scale.



Reliable:

Optional support for multiple availability zones. 99.95% service level agreement.



Client support:

Full support for the range of clients offered by Confluent, including Java, Python, C/C++, Go, .NET, as well as the Kafka Streams API.



Flexible:

Configurable retention period, storage, and throughput rate to suit your workload.



Kafka expertise:

Backed by the team that created Kafka and that has the most extensive experience operating it at scale.

* Google Cloud Platform and Microsoft Azure are slated for future releases.

Get Started in Minutes

Confluent Cloud provides an automation-friendly Command Line Interface (CLI) for use with the service. Using the CLI, developers can configure Kafka topics, run tests for producing and consuming messages, and query the overall health of the cluster and any messages flowing through it.

Confluent Cloud can be tuned across three core variables, as described below:

Throughput (write/read - MB/s)	Retention	Availability
0.5/1.0	7 days	Standard - 1 availability zone
2.0/4.0	14 days	Advanced - 3 availability zones with SLA
5.0/10.0	30 days	
<i>Custom throughput plans available</i>	<i>Custom retention and storage available for compacted topics</i>	

To learn more, visit www.confluent.io/confluent-cloud

Apache, Apache Kafka, Kafka, and associated open source project names are trademarks of the Apache Software Foundation