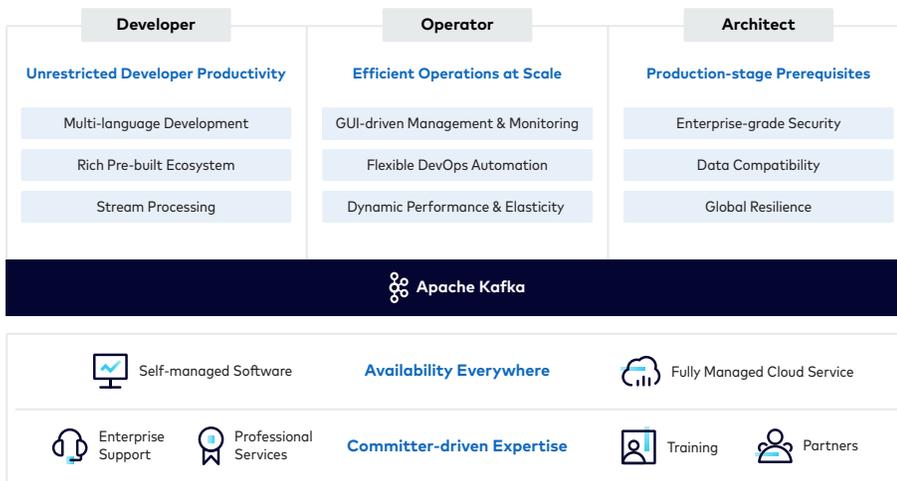


# Confluent Platform

Enterprise Platform for Data in Motion Built by the Original Creators of Apache Kafka®

Data in motion is a new paradigm that supports collecting a continuous flow of data throughout an organization and processing it in real-time. Originally developed at LinkedIn by the founders of Confluent, Apache Kafka is the industry standard infrastructure for setting data in motion and is used by organizations around the world to deliver exceptional customer experiences and data-driven, backend operations.

Confluent Platform is an enterprise-ready platform that completes Kafka with advanced capabilities designed to help accelerate application development and connectivity, enable transformations through stream processing, simplify enterprise operations at scale and meet stringent architectural requirements.



## What Customers are Saying

"Walmart is a \$500 billion in revenue company, so every second is worth millions of dollars. Having Confluent as our partner has been invaluable. We have expert support when we need it, best practices blueprints, and insight into new features to the product. Kafka and Confluent are the backbone of our digital omnichannel transformation at Walmart."

Chris Kasten  
Vice President | Walmart Cloud

"We love open source, but at the same time we're not a startup. We're a large financial institution that works with world-class organizations, and we need services that make it easier for us to sleep at night. Confluent Platform is very reliable; it's never down. It has become our backbone."

Kaspar Situmorang  
Executive Vice President | Bank Rakyat Indonesia

## Why Confluent Platform?

### Unrestricted Developer Productivity

Democratize Kafka for a wider range of developers and accelerate how fast they build real-time applications that leverage data in motion.

### Efficient Operations at Scale

Minimize operational complexity while ensuring high performance and scalability as data in motion grows through your organization.

### Production-stage Prerequisites

Architect the platform with the foundational enterprise-level attributes needed to implement data in motion use cases in production.

### Availability Everywhere

Deploy on-premises as either a cloud-native system on Kubernetes or on bare metal and virtual machines, or leverage a fully managed cloud service with [Confluent Cloud](#).

### Committer-driven Expertise

Leverage the world's foremost Kafka experts, who work directly with Kafka committers to assist you throughout the application development lifecycle.

# Features

	Apache Kafka	Confluent Platform
<b>Apache Kafka<sup>(1)</sup></b> with Connect and Streams	●	●
<b>Multi-language Development</b>		
<b>Clients<sup>(1)</sup></b> C++, Python, Go, and .NET		●
<b>REST Proxy<sup>(2)</sup></b> access to Kafka from any network connected device		●
<b>Admin REST APIs<sup>(2)</sup></b> RESTful interface for performing admin operations		●
<b>Rich Pre-built Ecosystem</b>		
<b>Connectors<sup>(1)(2)(3)</sup></b> over 100 supported, including S3, Elastic, HDFS, JDBC		●
<b>MQTT Proxy<sup>(3)</sup></b> access to Kafka from MQTT devices and gateways		●
<b>Schema Registry<sup>(2)</sup></b> central registry to ensure data compatibility		●
<b>Stream Processing</b>		
<b>Apache Flink<sup>(1)</sup></b> real-time stream processing for on-prem workloads with expert support		●
<b>ksqldb<sup>(2)</sup></b> framework that exposes Kafka Streams functionality through SQL		●
<b>GUI-driven Management and Monitoring</b>		
<b>Control Center<sup>(3)</sup></b> GUI to manage and monitor Kafka at scale		●
<b>Health+<sup>(3)</sup></b> intelligent alerts and cloud-based monitoring dashboards		●
<b>Flexible DevOps Automation</b>		
<b>Confluent for Kubernetes<sup>(3)</sup></b> complete, declarative API to deploy on Kubernetes		●
<b>Ansible<sup>(1)</sup></b> deployment automation on non-containerized environments		●
<b>Dynamic Performance and Elasticity</b>		
<b>Self-Balancing Clusters<sup>(3)</sup></b> automated partition re-balancing across brokers in the cluster		●
<b>Tiered Storage<sup>(3)</sup></b> offloading of older Kafka data to object storage while maintaining transparent access		●
<b>Enterprise-grade Security</b>		
<b>Role-Based Access Control<sup>(3)</sup></b> granular authorization of access by users/groups		●
<b>Structured Audit Logs<sup>(3)</sup></b> user actions logs stored in dedicated Kafka topics		●
<b>Secret Protection<sup>(3)</sup></b> encryption of sensitive information (e.g. passwords)		●
<b>Global Resilience</b>		
<b>Cluster Linking<sup>(3)</sup></b> highly available, consistent, and real-time bridge between on-prem and cloud environments		●
<b>Multi-Region Clusters<sup>(3)</sup></b> single Kafka cluster stretched across multiple data-centers with automated client failover		●
<b>Replicator<sup>(3)</sup></b> asynchronous replication based on the Kafka Connect framework		●
<b>Enterprise Support</b>		
24x7x365 from the Kafka experts		●

(1) Open Source (2) Community licensed feature (3) Commercially licensed feature

## Industry Recognition

- Top 10 Ranking on Forbes Cloud 100
- 2020 LinkedIn Top Startup
- Google Cloud Technology Partner of the Year Award
- JPMorgan Chase Hall of Innovation
- Morgan Stanley CTO Innovation Award
- Bank of America Recognition for Enterprise Technology Innovation

## Community Leaders

Confluent organizes the world's largest Kafka community event, [Kafka Summit](#), sponsors meetup groups and leads the [Kafka Slack community](#).

## Industry Experts

Confluent employs a very large team of [Apache Kafka committers](#) and experts, enabling us to ensure the success of your project through [enterprise support](#), [professional services](#) and [training](#).

## Get Started

Download [Confluent Platform](#) and try it free on a single broker or free for 30 days on unlimited brokers, or [sign up for Confluent Cloud](#) and start streaming in minutes.

Visit [confluent.io/get-started](https://confluent.io/get-started) for more details.