


Confluent Public Sector

Confluent Public Sector- knows government missions are driven by a constant stream of events, and that success depends on the ability to react as quickly as possible. Confluent offers an enterprise event streaming platform built by the original creators of Apache Kafka®. Confluent enables government agencies to utilize data as a continually updating stream of events, rather than discrete snapshots. Run your agency by building real-time applications with historical context - all based on a universal event pipeline.

The Confluent Platform Provides a Universal Event Pipeline Enabling Contextual Driven Apps



Confluent Platform Options:



Sensor Data Collection and Processing

Organizations across the government are leveraging Kafka to handle real-time feeds of sensor data. By publishing streams of events, they enable faster mission response, and multi-source data fusion. Disparate teams can deploy algorithms for detection and analysis and publish their results for others to consume and act on. Sensor data can span everything from telemetry to traditional SIGINT.



Data Ingest Pipeline

Kafka often serves as an ingestion pipeline and buffer for specific data sinks such as relational databases or ELK. Often times applications, databases, or indexes have trouble keeping up with ingestion rates particularly with spiky traffic patterns that can lead to dropped data. Kafka buffers incoming data until tools have a chance to catch up.



Legacy Messaging Modernization

The traditional notion of publish and subscribe, to decouple and enable asynchronous communication, remains a useful model for numerous applications. Legacy messaging technologies, however, have a number of limitations. With its modern, distributed architecture, Kafka delivers performance, massive scale, durability and availability. Kafka's ability to persist data also enables simplified data distribution and retroactive mission processing.



Digital Modernization of Citizen Engagement

Confluent's event streaming platform gives Government agencies the ability to interact more meaningfully with citizens over the course of their lifetime. Opportunities for improvements include a richer website experience, faster processing and vetting of applicants, rapid identification and action against fraud, faster enrollments or registration, and greater transparency. By accelerating these important functions, even performing them in real-time, Confluent Platform is designed to transform the nature of Citizen Engagement.



Efficient operations at scale

Minimize operational complexity while ensuring high performance and scalability as event streaming grows through your organization.



Freedom of choice

Deploy on-premises, in public or hybrid cloud, from bare-metal to Kubernetes, 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.



Central Nervous System

Achieving the unfulfilled vision of the Enterprise Service Bus (ESB) is now a reality with modern technical architecture. Enterprise data flows through Kafka from sources of truth to the many downstream consumers. This eliminates the complexity of point to point integration. Lyft, Uber, LinkedIn, Netflix and many other organizations deploy and manage Kafka clusters to provide access to all the real-time data that a digital organization needs.



Cyber Data Collection and Dissemination

This is now one of the fastest growing usage patterns with Confluent - Cyber Data (network events, log events, etc.) flow into Kafka, which then feeds into one or more tools like Splunk, ArcSight, or Elastic, Logstash and Kibana (ELK). In addition, real-time analytics, processing, and alerting takes place directly against the streaming data. This brings countless advantages such as allowing usage of best of breed tools, cost reduction by stratifying data and the tools used for them (a subset into Splunk but a large amount in ELK), and enabling a hub/spoke distributed method of handling data.



Data Distribution and Replication

Confluent replication, built on top of Kafka, gives a reliable and flexible way to replicate data thereby ensuring high data fidelity. This can be hub and spoke, where there is edge data collection/processing and a central aggregation. It could be simple data center to data center replication. It's particularly useful in environments with multiple applications or tools that require the data.



Hybrid Cloud and Replatforming

Kafka provides an on-ramp to migrate to a modern architecture by building a persistent bridge between data center and cloud with Confluent Replicator. Data can flow into Kafka from legacy systems and be distributed to numerous new applications. Legacy and new platforms can co-exist until the time the legacy system is taken offline. This is a common pattern for mainframe customers.



Unrestricted developer productivity

Democratize Kafka for a wider range of developers and accelerate how fast they build event streaming applications.



Production-stage prerequisites


Architect the platform with the foundational enterprise-level attributes needed to implement event streaming in production. Confluent makes it easier to develop, deploy, secure and operate at enterprise scale.

CONFLUENT CLOUD

Cloud-native service for data in motion built by the original creators of Apache Kafka®

Today's consumers have the world at their fingertips and hold an unforgiving expectation for end-to-end real-time brand experiences. Data in motion is the underlying, fundamental ingredient to any truly connected customer experience. It provides a continuous supply of real-time event streams coupled with real-time stream processing to power the data-driven backend operations and rich front-end experiences necessary for any business to succeed within today's competitive, consumer-driven markets.

Set your data in motion while avoiding the headaches of infrastructure management and focus on what matters most: your business. Built by the original creators of Apache Kafka, Confluent Cloud is a fully managed, cloud-native service for connecting and processing all of your real-time data, everywhere it's needed.



Build and launch real-time apps faster

Unless you're in the business of managing data infrastructure, self-managing open source Kafka is a challenging and costly path to take. When accounting for cluster sizing, scaling, Zookeeper management, integrations development, patches, upgrades, and more, the cost of managing the open source technology grows over time and seriously limits your ability to focus on innovation and building true value. Leverage a service that allows you to fully focus on your core business priorities. Confluent provides a truly fully managed Kafka service that is cloud-native, complete, and available everywhere you need it.

Cloud Native

With Confluent Cloud, you can provision serverless Kafka clusters on demand that scale elastically between 0-100 MBps or get you to GBps+ scale with just a few clicks. Instant elasticity like this means you can easily scale up to meet unexpected demand and just as quickly scale back down to manage cost. You pay for what you need, when you need it, and no more.

You can also store unlimited amounts of data on your Kafka clusters, without any upfront capacity planning or provisioning. Kafka becomes a system of record so you can do more with your real-time events. Most importantly, leveraging a cloud-native platform means running with the confidence that it just works. Our highly available clusters are backed by a 99.95% SLA and always on the latest stable version of Kafka—upgraded and patched in the background like a true cloud-native service.

Complete

Confluent provides way more than just Kafka. You get a complete platform for data in motion allowing you to build real-time apps quickly, reliably, and securely. We provide out-of-the-box, fully managed connectors for the most popular data sources/sinks in the Kafka ecosystem, the industry's only governance suite for data in motion, stream processing with a streaming database, and more—all fully managed in the same cloud UI.

There's no cutting corners when it comes to data safety and compliance. Confluent secures your Kafka environment with encryption at rest and in transit, role-based access controls, Kafka ACLs, and SAML/SSO for authentication. More granular control with private networking and BYOK for your Kafka environments is available as well. Finally, we design with compliance and privacy in mind with SOC 1/2/3 certifications, ISO 27001 certification, HIPAA/GDPR/CCPA readiness, and more.

Everywhere

Studies have shown that 80 percent of companies use more than one cloud service provider. If you're one of the lucky ones with a single provider, that's great. If not, a Kafka deployment that only works in a single environment undercuts your ability to leverage data in motion across your entire business and every customer experience. Confluent lets you link Kafka clusters that sync in real time, so your events are available anywhere—across multiple public or private clouds.

Features

Cloud Native

Run at massive scale cost-effectively without the operational burden

Elastic

- Self-serve provisioning of serverless Kafka clusters with no complex sizing requirements
- Elastic scaling between 0-100 MBps and GBps+ scale with just a few clicks
- Pay only for what you actually use

Infinite

- Infinite data storage and retention within Kafka topics (AWS & Google Cloud)

Highly Available

- Guaranteed 99.95% uptime SLA with multi-AZ replication
- Zero downtime Kafka upgrades & bug fixes

Complete

Go way beyond Kafka to build real-time apps quickly, reliably, and securely

120+ Connectors

- Instantly integrate clusters to the most popular data sources and sinks within the Kafka ecosystem including MongoDB, Elastic, Snowflake, and more

Stream Governance

- Discover, understand, and trust your data streams with the industry's only governance suite for data in motion including fully managed Schema Registry & Validation

Streaming Database

- Launch real-time and continuous stream processing apps with ksqlDB's lightweight SQL syntax

Enterprise-grade Security

- Build quickly while upholding strict security and compliance requirements with data encryption at rest and in transit, BYOK encryption, RBAC, Kafka ACLs, SAML/SSO authentication, and private networking
- SOC 1/2/3 certifications, ISO 27001 certification, HIPAA/GDPR/CCPA readiness, and more

Everywhere

Work with real-time events in the cloud, across clouds, or on-premises

Public Cloud & Multi-Cloud

- Confluent Cloud is provided fully managed on AWS, Azure, and Google Cloud



Cluster Linking

- Simplify geo-replication and multi-cloud data movement with fully managed cluster linking and perfectly mirrored offsets between global clusters

Deployment Flexibility

- Build a persistent bridge from on-premises to cloud with [Confluent Platform](#) and a hybrid Kafka service



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 projects through [enterprise support](#), [professional services](#), and [training](#).



Industry Recognition

- [Top 10 ranking on Forbes Cloud 100](#)
- [2020 LinkedIn Top Startup](#)
- [3x Google Cloud Technology Partner of the Year](#)
- [JPMorgan Chase Hall of Innovation](#)
- [Morgan Stanley CTO Innovation Award](#)
- [Bank of America recognition for Enterprise Technology Innovation](#)



Get started with Confluent Cloud for free and start working with real-time data in minutes.

New signups receive up to \$400 to spend on Confluent Cloud during their first 60 days.

Signup is also available through cloud marketplaces: [AWS](#), [Azure](#), and [Google Cloud](#).

Just getting started with Kafka and looking to learn?

Visiting [Confluent Developer](#) is the easiest way to learn Apache Kafka and event streaming.

What customers are saying



"We chose an event-driven architecture as the core of our platform, for which we needed a messaging service that gave us all the guarantees—it had to be extremely scalable, highly available, and simple to use. Kafka hit all of these markers, and by using Confluent Cloud, our team was able to reduce the bottom line and operational burden."

Ravi Vankamamidi | Senior Director, Technology, at Expedia Group



"We truly think of Confluent Cloud and Kafka as the central nervous system of our business, spanning everything from the customer-facing applications to distribution center operations from a technology perspective."


Chirag Dadia | Director of Engineering at Nuuly

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