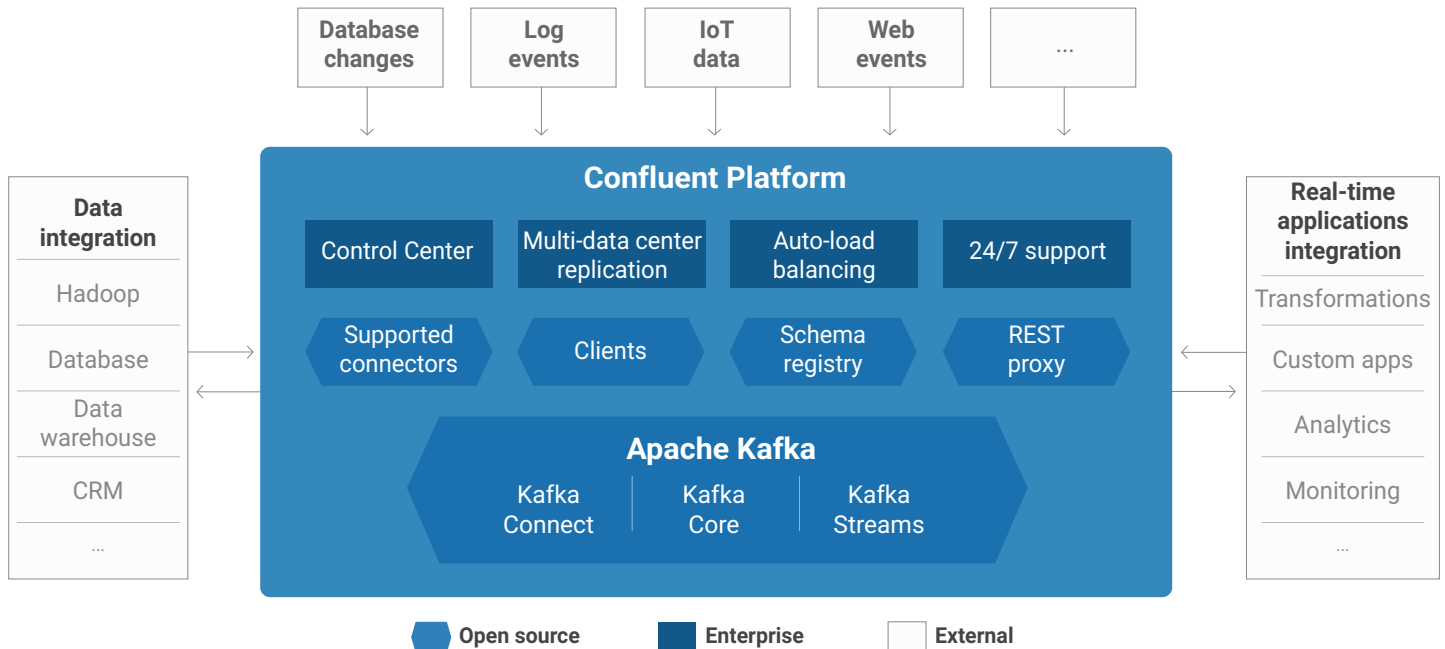


Monitoring, Administration and Operations for Running Apache Kafka at Scale

Organizations around the world rely on Apache Kafka™ to run some of the largest data streaming pipelines, to build microservices applications, and to support IoT initiatives. Confluent Enterprise is the only enterprise stream platform built entirely on Kafka that makes implementing and managing an enterprise streaming platform with Kafka easy, reliable, and secure.



At the core of Confluent Enterprise is Apache Kafka, a modern distributed system for moving data in real time. It allows a single cluster of computers to perform all the transport between an organization’s data systems regardless of their number or size. Unlike other solutions that leverage only a part of Kafka to serve as a fast pipe for loading data, Confluent supports the full range of Kafka’s capabilities including the Connect and Streams APIs for building robust, mission critical real-time data solutions.

Built around 100% open source Apache Kafka, Confluent Enterprise provides a rich set of streaming capabilities free from lock-in. Because Kafka is open source, data flows through it commitment free. With a broad and expanding partner ecosystem, Confluent integrates with many existing systems. Pre-built connectors for leading tools such as MongoDB, SAP and Hadoop make it simple to build the ideal streaming solution.

Confluent Enterprise is designed to meet the stringent security and governance requirements common in most industries today. With built-in security for data stored in or streaming through the platform, and data consistency capabilities to ensure any data flowing through the streaming platform is adequately governed, it’s possible to build trusted streaming data solutions for even the most regulated enterprise.

Finally, Confluent Enterprise delivers the operations, monitoring, and administration tools for running Kafka at scale. With capabilities for simplifying replication across datacenters, auto rebalancing to keep clusters performing efficiently, and an operational Control Center for ensuring everything is running smoothly, only Confluent offers the confidence needed to architect a complete streaming solution with Kafka.

What is a Streaming Platform?

A streaming platform has three key characteristics:



Publish and subscribe to streams of data



Store streams of data within the platform



Process streams of data without requiring an additional cluster

A streaming platform serves as a “central nervous system” for data, connecting multiple disparate data sources and applications for building mission critical real-time applications. The unique capabilities of a complete streaming platform enable several key use cases, including:



Real-time ETL pipelines

Real-time pipelines are pretty much the same in all industries and companies—moving, transforming, and integrating data is a well-defined technical problem regardless of the business it happens in. Even different varieties of pipeline really aren't all that different—pipelines to Hadoop clusters or between on premise and cloud or for real-time analytics really only vary by volume and latency. Consumers and flow any amount of data through the platform.



Streaming applications

Real-time applications are very specific to the industry and company: health-care companies want to monitor blood glucose in real-time, finance companies want to process market data, retailers want to manage inventory. Organizations across industry are building complete streaming applications using Confluent.



Microservices

Microservices are often built upon a streaming platform. Instead of a large monolithic application, microservices deliver small, decoupled processes that execute a constrained feature set against a stream. These are ideal for a streaming platform because they work on real-time streams, depend on reliable message delivery, and might not warrant the cost of deploying a separate stream processing framework.

The Benefits of Kafka



All Your Data

Kafka's architecture makes it an ideal centralized service for all the data in your organization. Get all your data from disparate systems flowing through Kafka — unlocking your data to fuel your growth.



Highly Scalable

Kafka scales to trillions of messages a day and thousands of producers and consumers without loss of performance, allowing you to build a real-time streaming platform connecting all of your infrastructure.



Real-Time

Kafka satisfies the millisecond latency requirements of streaming applications, is scalable enough to handle very high volume log and event data, and fault-tolerant enough for critical data delivery. Kafka allows you to run very large data streams in real time across your organization.

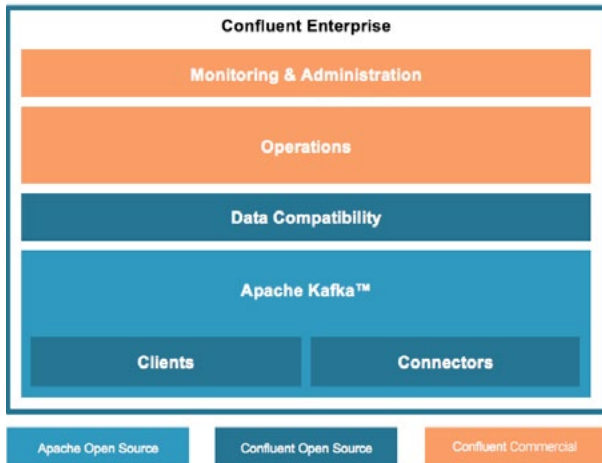


Secure

Kafka includes the encryption, authentication, and authorization capabilities needed to ensure sensitive data is protected and accessible only by authorized parties.

Features

Confluent enables enterprises of any size to integrate real-time streams into their business processes. Only Confluent Enterprise includes monitoring, administration, operations, and data compatibility capabilities required to run Kafka at scale.



7 of the top 10 global banks, 7 of the top 10 insurance companies, and 9 of the top 10 telecommunications providers rely on Apache Kafka to run their businesses. Confluent extends the value of Kafka by adding the enterprise capabilities to run Kafka for more mission critical services.



Clients and Connectors

Confluent Enterprise includes the most producer and consumer clients of any Kafka based product. It includes clients for Java, C/C++, Python, .NET and several others for working with a broader set of applications.

Confluent also includes a wide array of connectors making it easy to connect to existing systems including databases, Hadoop, Elasticsearch, or Amazon S3.

A REST Proxy provides access to Kafka from any network connected device that uses HTTP. Additionally, it provides a way for clients to connect to a Kafka cluster through web proxies or other intermediate services.



Data Compatibility

Confluent Schema Registry provides a central registry for the format of the data in each Kafka topic and provides a central service that helps to make changes to data formats easy and backwards compatible. Confluent Schema Registry supports the Apache Avro data format, and provides a RESTful interface to manage schemas and schema changes.



Operations

With Auto Data Balancer, Confluent improves resource utilization and reliability by automatically balancing data between Kafka servers in a cluster. Users experience a lower probability of failures, improved performance and a streamlined operations team.

In a multi-cluster environment, Replicator makes it easy to set up replication and maintain clusters in many geographical locations with centralized configuration. Clusters in each datacenter are synchronized so applications and configurations in separate datacenters are in sync.



Monitoring and Administration

Confluent Control Center provides a comprehensive management system for Kafka, enabling the visibility and operational strength needed to manage a Kafka environment. Control Center delivers real-time analysis of the performance of Kafka, allowing data teams to drill into topics, producers, consumers, and more to understand what's happening with their data pipelines.

Commitment to Kafka

Confluent was founded by the engineering team at LinkedIn that originally created Apache Kafka, employs many of the active committers to the open source project and continues to invest in a thriving Kafka community. You will not find better resources to support your Kafka implementation than those that Confluent can provide. The Confluent team includes the foremost experts in stream data management of all types, particularly Kafka.



Enterprise Support

Confluent’s engineering and support teams include the largest team of contributors to Apache Kafka, with many years of development and operational experience with distributed systems. Confluent provides 24/7 enterprise support, critical bug fixes, indemnification, training, and operational advice to customers.



Professional Services

With a range of tailored expert consulting engagements that align with each stage of a project from inception through to post-production launch, Confluent specializes in ensuring deployment success. We offer engagements such as Application Architecture, AWS Cloud Deployment and Health Check, and also design custom engagements to meet specific needs.



Training

Developers and operators often need help in getting up to speed on Kafka, which is why we offer in depth training classes to ensure everyone has the skills they need to run a successful streaming platform deployment.



Get Started with Confluent Enterprise

Getting started with Confluent is quick and easy. The platform is available commitment free via an open source download. Enterprise capabilities are available through our enterprise subscription or via a 30-day trial. To learn more, visit the download center at www.confluent.io/download.

Feature	Description	Apache Kafka	Confluent Open Source	Confluent Enterprise
Apache Kafka	High throughput, low latency, high availability, secure, distributed streaming platform	●	●	●
Kafka Connect API	Advanced API connects external sources/destinations into Kafka	●	●	●
Kafka Streams API	Simple library that enables streaming application development within the Kafka framework	●	●	●
Additional Clients	Supports non-Java clients; C, C++, Python, etc.		●	●
REST Proxy	Provides universal access to Kafka from any network connected device via HTTP		●	●
Schema Registry	Central registry for the format of Kafka data – guarantees all data is always consumable		●	●
Pre-Built Connectors	HDFS, JDBC, elasticsearch and other connectors certified and supported by Confluent		●	●
JMS Client	Support for legacy Java Message Service (JMS) applications consuming and producing directly from Kafka			●
Confluent Control Center	Enables easy connector management and stream monitoring			●
Auto-data Balancing	Rebalancing data across clusters to remove bottlenecks			●
Multi-Data Center Replication	Multi-datacenter replication simplifies and automates multi-datacenter Kafka clusters			●
Support	Enterprise class support to keep your Kafka environment running at top performance	Community	Community	24x7x365