

Confluent Cloud for Apache Flink®





Simple, Serverless Stream Processing

Stream processing acts as the compute layer for Apache Kafka®, enabling teams to combine and reshape their data streams in-flight to power real-time applications and pipelines. Apache Flink® is the stream processor of choice for many of the world’s largest real-time systems and is used by innovative companies like AirBnB, Uber, Netflix, and TikTok to support mission-critical streaming workloads.

With Confluent Cloud for Apache Flink, customers can easily build high-quality, reusable data streams with the industry’s only cloud-native, serverless Apache Flink service, fully integrated with Apache Kafka on Confluent Cloud across all three major clouds.

In addition, Confluent Platform for Apache Flink®, currently in Limited Availability, allows you easily leverage stream processing for on-prem or private cloud workloads with support from the Flink experts.

 Effortlessly filter, join, and enrich your data streams with Flink, the de facto standard for stream processing

 Enable high-performance and efficient stream processing at any scale, without any operational burden

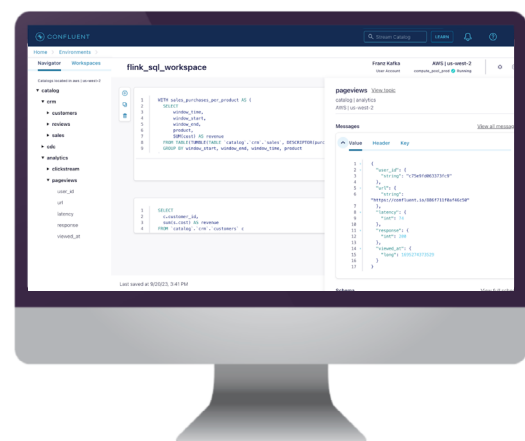
 Experience Apache Kafka and Apache Flink seamlessly integrated as a complete, unified platform

Why Confluent Cloud for Apache Flink?

Data streaming reimagines how data is made available within an organization, continuously ingesting, processing, governing, and sharing data as a product that is instantly available, usable, and trustworthy everywhere. Stream processing plays a critical role in the infrastructure stack for data streaming by enabling developers to filter, join, aggregate, and transform data on-the-fly, making it more usable for downstream systems and applications.

Traditional batch-oriented processing downstream of Kafka cannot achieve the same outcomes. It leads to applications built on stale and inconsistent data, requiring expensive, inefficient, and redundant processing to clean and enrich data for each downstream system and application. Although Apache Flink is widely recognized for its performance and rich feature set among stream processing frameworks, self-managing Flink (like Kafka) can be challenging due to operational complexity, a steep learning curve, and high costs for in-house support.

Confluent Cloud for Apache Flink offers a cloud-native, serverless Flink service for processing data in-flight and creating high-quality, reusable streams in real time. Effortlessly filter, join, and enrich data streams with Flink, the industry-leading stream processing framework, at any scale with zero ops burden, and experience Kafka and Flink as a unified platform with integrated monitoring, security, and governance.



Effortlessly filter, join, and enrich your data streams with Flink, the de facto standard for stream processing

Real-time Processing

Power low-latency applications and pipelines that react to real-time events and provide timely insights

Data Reusability

Share consistent and reusable data streams widely with downstream applications and systems

Data Enrichment

Curate, filter, and augment data on-the-fly with additional context to improve completeness, accuracy, and compliance

Efficiency

Improve resource utilization and cost-effectiveness by avoiding redundant processing across silos

Enable high-performance and efficient stream processing at any scale, without any operational burden

Fully Managed

Easily develop Flink applications with a serverless, SaaS-based experience instantly available and without ops burden

Elastic Scalability

Automatically scale up or down to meet the demands of the most complex workloads without overprovisioning

Usage-based Billing

Pay only for resources used instead of infrastructure provisioned, with scale-to-zero pricing

Continuous, No-touch Updates

Build on an always up-to-date platform with declarative, versionless APIs and interfaces, backed by a 99.99% SLA

Experience Apache Kafka and Apache Flink seamlessly integrated as a complete, unified platform

Enterprise-grade Security

Secure stream processing with built-in identity and access management, RBAC, and audit logs

Stream Governance

Enforce data policies and avoid metadata duplication leveraging native integration with Stream Governance

Monitoring

Ensure the health and uptime of your Flink queries in the Confluent UI or via third-party monitoring services

Connectors

Connect to and from any app & system with **70+ pre-built, fully managed connectors.**



“Stream processing is critical for identifying and protecting against security risk in real time. With Confluent’s fully managed Flink offering, we can access, aggregate, and enrich data from IoT sensors, smart cameras, and Wi-Fi analytics, to swiftly take action on potential threats in real time, such as intrusion detection. This enables us to process sensor data as soon as the events occur, allowing for faster detection and response to security incidents without any added operational burden.”

Vinay Krishna Patnana | Engineering Manager at Cisco Meraki