



Confluent Unveils Apache Flink® on Confluent Cloud, Making it Easier to Build Real-Time Applications with Stream Processing on a Unified Platform

September 26, 2023

Confluent launches the industry's only serverless, cloud-native Flink service to simplify building high-quality, reusable data streams

Confluent expands Stream Governance capabilities with Data Portal, so teams can easily find all the real-time data streams in an organization

New Confluent Cloud Enterprise offering lowers the cost of private networking and storage for Apache Kafka

SAN JOSE, Calif.--(BUSINESS WIRE)--Sep. 26, 2023-- [Confluent, Inc.](#) (NASDAQ: CFLT), the data streaming pioneer, today announced the open preview of [Apache Flink® on Confluent Cloud](#), a fully managed service for stream processing that makes it easier for companies to filter, join, and enrich data streams with Flink. With data streams processed on the fly, organizations are able to react and innovate faster with timely insights. In addition, Confluent announced Data Portal to help teams discover all the real-time data streams within their organizations, Enterprise Clusters to cut [Apache Kafka®](#) costs for serverless private networking, and reduced storage costs for Confluent Cloud.

Data streaming is a critical business requirement – [72%](#) of IT leaders use it to power mission-critical systems, according to the [2023 Data Streaming report](#). But streaming data is only part of the puzzle. To derive the full value of data streams, organizations need to combine, enrich, and reshape those streams with other data from across their organization. In a data streaming platform, stream processing acts as the compute layer that turns data into valuable insights and action.

As one of the top projects of the Apache Software Foundation, Flink has emerged as the de facto standard for stream processing. Flink's high performance, rich feature set, and robust developer community make it one of the most popular choices for large-scale, high throughput, and low-latency stream processing. Since Kafka has become the standard technology for [data streaming](#), Flink is often used in tandem with Kafka to support companies' mission-critical workloads.

"Companies globally are leveraging data streaming with Apache Kafka to power real-time experiences," said Shaun Clowes, Chief Product Officer at Confluent. "While data streaming connects data across an organization, stream processing tools like Apache Flink make it possible to act on that data in real time and accelerate the development of new use cases. Our managed Kafka and Flink offerings are like the peanut butter and jelly of data streaming. They work better together on a unified platform to connect and enrich data so teams can uncover more insights and deliver impactful experiences that move the needle."

Flink makes it possible for ride sharing companies to match drivers and riders, for banks to alert users on fraudulent activity, and for social media companies to offer personalized content recommendations.

"Stream processing is critical for identifying and protecting against security risks in real time," said Vinay Krishna Patnana, Engineering Manager at Cisco Meraki. "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."

However, like Kafka and other open source technologies, self-managing Flink is operationally complex, has a steep learning curve, and comes with expensive infrastructure and management costs.

Confluent's Apache Flink service simplifies stream processing, enabling faster application development

[Apache Flink](#) on Confluent Cloud allows teams to create high-quality, reusable data streams that can be delivered anywhere in real time. With Confluent's fully managed and elastically scalable Flink service, teams can reduce the architectural complexity and operational burdens of stream processing to enrich streaming data and power innovative use cases.

"Stream processing involves real-time integration and transformation of streams of event data, powering systems across industries ranging from fraud detection in finance to personalized recommendations in e-commerce," said Matt Aslett, Vice President and Research Director at Ventana Research. "When used in combination, Apache Flink and Apache Kafka can enable data reusability and avoid redundant downstream processing. The delivery of Flink and Kafka as fully managed services delivers stream processing without the complexities of infrastructure management, enabling teams to focus on building real-time streaming applications and pipelines that differentiate the business."

Teams using Apache Flink on Confluent Cloud can:

- **Effortlessly filter, join, and enrich data streams with Flink**, the standard for stream processing.
- **Enable high-performance and efficient stream processing at any scale**, without the complexities of infrastructure management.
- **Experience Kafka and Flink as a unified platform**, with fully integrated monitoring, security, and governance.

Apache Flink is available as an open preview for [Confluent Cloud](#) customers using AWS in select regions for testing and experimentation purposes.

General availability is coming soon.

Additional Confluent Cloud innovations

Data Portal increases developer productivity with a simple interface for discovering, accessing, and leveraging real-time data streams

With open source Kafka and other managed Kafka solutions, developers often run into problems finding and accessing the relevant data streams they need to build real-time applications and pipelines. This fragmented process wastes valuable development cycles, restricts productivity, and slows innovation.

To address these challenges, Data Portal, an expansion of the [Stream Governance](#) suite, leverages Confluent's Stream Catalog capabilities to simplify the developer experience. With its flexible and self-service interface, Data Portal gives teams a secure way to find and access all the data streams flowing throughout their organization, speeding up the development of real-time applications and products.

Data Portal enables teams to:

- **Search, discover, and explore existing topics, tags, and metadata** across the organization with end-to-end visibility to choose the data most relevant for their projects.
- **Seamlessly and securely request access to data streams** and trigger an approval workflow that connects the user with the data owner, all within the Confluent Cloud UI.
- **Easily build and manage data products to power streaming pipelines and applications** by understanding, accessing, and enriching existing data streams.

Confluent Cloud customers will be able to experience Data Portal's features to easily find data streams soon.

Kora drives new efficiencies and cost savings

[Kora](#), Confluent's cloud-native Apache Kafka engine, enables Confluent to efficiently manage tens of thousands of Confluent Cloud clusters, resulting in cost savings for our customers.

Enterprise Clusters maximize efficiency on private networks while minimizing operational challenges

Confluent is expanding its serverless capabilities to those who need private networking with [Enterprise clusters](#) on Confluent Cloud. Built on the Kora Engine and developed for teams with stringent security and networking requirements, teams using Enterprise clusters can establish secure and direct communication between VPCs and Confluent Cloud without exposing their data to the internet.

With Enterprise Clusters on Confluent Cloud, teams can:

- **Easily and securely connect private environments to Confluent Cloud** with simplified, reusable setup and secure network isolation.
- **Optimize resource and cost efficiency with auto-scaling clusters** to meet any demand.
- **Eliminate manual sizing, provisioning and ongoing management** with automated operations and intelligent data tiering powered by our Kora Engine.

Enterprise clusters are now available on AWS PrivateLink to securely access resources in Confluent Cloud and simplify network administration.

Reduced storage pricing for Confluent Cloud enables more data storage for less

Confluent Cloud now offers storage for [20%](#) less, effective October 1, 2023. Teams can still retain all real-time and historical events without limits, helping power more streaming use cases, including event sourcing, artificial intelligence, and stream processing, at a more affordable price. [Try Confluent Cloud for free](#) in minutes.

These innovations are made possible by the Kora Engine and underscore the benefits of Confluent's cloud-native platform. Over time, the new enterprise clusters and lower storage pricing are expected to drive lower TCO and better ROI, while improving the efficiency at which Confluent can deliver its complete data streaming platform to new and existing customers.

Learn more at Current!

These new features and capabilities will be demoed onstage at [Current](#), the data streaming industry event, taking place on September 26-27. Tune in today, September 26 at 8:45 am PT, to watch a keynote presentation on the evolution and impact of data streaming platforms. [Register here](#) to watch the keynote program and select sessions virtually or attend in person to see live demos at the Confluent booth. A live webcast and a replay of each presentation will be available on Confluent's website [here](#).

Additional Resources

- Explore all the new innovations of the [Confluent Q3 '23 Launch in this blog post](#)
- Learn more about [Confluent's Managed Flink Service](#)
- For deep dives into each feature, register for the upcoming [Confluent Q3 '23 Launch webinar](#)
- See how Confluent is [helping customers transform their business](#)

About Confluent

Confluent is the data streaming platform that is pioneering a fundamentally new category of data infrastructure that sets data in motion. Confluent's

cloud-native offering is the foundational platform for data in motion—designed to be the intelligent connective tissue enabling real-time data, from multiple sources, to constantly stream across the organization. With Confluent, organizations can meet the new business imperative of delivering rich, digital front-end customer experiences and transitioning to sophisticated, real-time, software-driven back-end operations. To learn more, please visit www.confluent.io.

Confluent and associated marks are trademarks or registered trademarks of Confluent, Inc.

Apache® and Apache Kafka® are either registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries. No endorsement by the Apache Software Foundation is implied by the use of these marks. All other trademarks are the property of their respective owners.

This press release contains forward-looking statements, including among other things, statements regarding the benefits, performance, features and use cases of Apache Flink on Confluent Cloud, the role of stream processing for organizations, market acceptance of data streaming, synergies and benefits of managed Kafka and Flink, benefits and features as well as roll-out timing of Data Portal and Kora, the anticipated benefits of new enterprise clusters and lower storage pricing including long-term benefits to the efficiency of Confluent's business, the conversion of open source users, and the adoption of Confluent's data streaming platform, and timing of certain pricing changes for the storage feature of Confluent Cloud. The words "believe," "may," "will," "ahead," "estimate," "continue," "anticipate," "intend," "expect," "seek," "plan," "project," and similar expressions are intended to identify forward-looking statements. These forward-looking statements are subject to risks, uncertainties, and assumptions. If the risks materialize or assumptions prove incorrect, actual results could differ materially from the results implied by these forward-looking statements. Except as required by law, Confluent assumes no obligation to and does not currently intend to, update any such forward-looking statements after the date of this release.

View source version on [businesswire.com](https://www.businesswire.com/news/home/20230926152351/en/): <https://www.businesswire.com/news/home/20230926152351/en/>

Media Contact

Natalie Mangan
pr@confluent.io

Source: Confluent, Inc.