



Confluent Launches New Connectors and On-Demand Scaling to Break Down Data Silos and Meet the Unpredictable Needs of Modern Business

January 19, 2022

Additions to the over 50 expert-built, fully managed connectors quickly modernize applications with real-time data pipelines

New controls to expand and shrink GBps+ cluster capacity enhance elasticity for dynamic, real-time business demands

New Schema Linking ensures trusted, compatible data streams across cloud and hybrid environments around the globe

MOUNTAIN VIEW, Calif.--(BUSINESS WIRE)--Jan. 19, 2022-- [Confluent, Inc.](#) (NASDAQ: CFLT), the platform to set data in motion, today announced the Confluent Q1 '22 Launch, which includes new additions to the industry's largest portfolio of fully managed data streaming connectors, new controls for cost-effectively scaling massive-throughput Apache Kafka® clusters, and a new feature to help maintain trusted data quality across global environments. These innovations help enable simple, scalable, and reliable data streaming across the business, so any organization can deliver the real-time operations and customer experiences needed to succeed in a digital-first world.

"Because of how we now consume data, companies and customers have a growing expectation of immediate, relevant experiences and communication," said Amy Machado, Research Manager, Streaming Data Pipeline, IDC in IDC's Worldwide Continuous Analytics Software Forecast, 2021–2025. "Real-time datastreams and processing will become the norm, not the exception."

However, for many organizations, real-time data remains out of reach. Data lives in silos, trapped within different systems and applications because integrations take months to build and significant resources to manage. In addition, adapting streaming capacity to meet constantly changing business needs is a complex process that can result in excessive infrastructure spend. Lastly, ensuring data quality and compliance on a global scale is a complicated technical feat, typically requiring close coordination across teams of Kafka experts.

"The real-time operations and experiences that set organizations apart in today's economy require pervasive data in motion," said Ganesh Srinivasan, Chief Product Officer, Confluent. "In an effort to help any organization set their data in motion, we've built the easiest way to connect data streams across critical business applications and systems, ensure they can scale quickly to meet immediate business needs, and maintain trust in their data quality on a global scale."

Introducing new additions to Confluent's cloud-native data streaming platform

With these latest innovations that are now generally available, Confluent continues to deliver on its vision of providing customers with a data streaming platform that is complete, cloud native, and everywhere.

Complete: Additions to the over 50 expert-built, fully managed connectors quickly modernize applications with real-time data pipelines

Confluent's newest connectors include [Azure Synapse Analytics](#), [Amazon DynamoDB](#), [Databricks Delta Lake](#), [Google BigTable](#), and [Redis](#) for increased coverage of popular data sources and destinations.

"Running the largest online marketplace for independent creators requires data in motion to better serve our users," said Joe Burns, CIO, TeePublic. "We are a small team tasked with making real-time user interaction data available in our data warehouse and data lake for immediate analysis and insights. By using fully managed Amazon S3 sink, Elasticsearch sink, Salesforce CDC source, and Snowflake sink connectors, we were able to quickly and easily build high-performance streaming data pipelines that connect our business through Confluent Cloud without any operational burden, accelerating our overall project timeline."

Available only on Confluent Cloud, Confluent's portfolio of [over 50 fully managed connectors](#) helps organizations build powerful streaming applications and improve data portability. These connectors, designed with Confluent's deep Kafka expertise, provide organizations an easy path to modernizing data warehouses, databases, and data lakes with real-time data pipelines:

- Data warehouse connectors: Snowflake, Google BigQuery, Azure Synapse Analytics, Amazon Redshift
- Database connectors: MongoDB Atlas, PostgreSQL, MySQL, Microsoft SQL Server, Azure Cosmos DB, Amazon DynamoDB, Oracle Database, Redis, Google BigTable
- Data lake connectors: Amazon S3, Google Cloud Storage, Azure Blob Storage, Azure Data Lake Storage Gen 2, Databricks Delta Lake

To simplify real-time visibility into the health of applications and systems, Confluent announced first-class integrations with [Datadog](#) and [Prometheus](#). With a few clicks, operators have deeper, end-to-end visibility into Confluent Cloud within the monitoring tools they already use. This provides an easier means to identify, resolve, and avoid any issues that may occur while returning valuable time for everything else their jobs demand.

"Delivering a best-in-class experience for homeowners and housing professionals around the world requires a holistic view of our business," said Mustapha Benosmane, Product leader, ADEO. "Confluent's integration with Datadog quickly syncs real-time data streams with our monitoring tool of choice with no operational complexities or middleware required. Our teams now have visibility into the health of all of our systems for reliable,

always-on services.”

Cloud Native: New controls to expand and shrink GBps+ cluster capacity enhance elasticity for dynamic, real-time business demands

To ensure services always remain available, many companies are forced to over-provision capacity for their Kafka clusters, paying a steep price for excess infrastructure that often goes unused. Confluent solves this common problem with [Dedicated clusters](#) that can be provisioned on demand with just a few clicks and include self-service controls for both [adding](#) and [removing capacity](#) to the scale of GBps+ throughput. Capacity is easy to adjust at any time through the Confluent Cloud UI, CLI, or API. With automatic data balancing, these clusters constantly optimize data placement to balance load with no additional effort. Additionally, minimum capacity safeguards protect clusters from being shrunk to a point below what is necessary to support active traffic.

“Ensuring our e-commerce platform is always up and available is a complex, cross-functional, and expensive effort, especially considering the major fluctuations in online traffic we see throughout the year,” said Cem Küççük, Senior Manager, Product Engineering, Hepsiburada (NASDAQ:HEPS). “Our teams are challenged to deliver the exact capacity we need at any given time without over-provisioning expensive infrastructure. With a self-service means to both expand and shrink cloud-native Apache Kafka clusters, Confluent allows us to deliver a real-time experience for every customer with operational simplicity and cost efficiency.”

Paired with Confluent’s new [Load Metric API](#), organizations can make informed decisions on when to expand and when to shrink capacity with a real-time view into utilization of their clusters. With this new level of elastic scalability, businesses can run their highest throughput workloads with high availability, operational simplicity, and cost efficiency.

Everywhere: New Schema Linking ensures trusted, compatible data streams across cloud and hybrid environments around the globe

“As enterprises begin to adopt event streaming more broadly, sharing event data is both more important and more common,” according to Maureen Fleming, IDC program VP, Intelligent Process Automation. “Capabilities like schema linking enable faster adoption, lower costs, and more trust in leveraging the data flowing across an enterprise.”

Global data quality controls are critical for maintaining a highly compatible Kafka deployment fit for long term, standardized use across the organization. With [Schema Linking](#), businesses now have a simple way to maintain trusted data streams across cloud and hybrid environments with shared schemas that sync in real time. Paired with [Cluster Linking](#), schemas are shared everywhere they’re needed, providing an easy means of maintaining high data integrity while deploying use cases including global data sharing, cluster migrations, and preparations for real-time failover in the event of disaster recovery.

Additional Resources

- For more details on the Confluent Q1 ‘22 Launch, check out this blog post: <https://www.confluent.io/blog/real-time-data-integrations-at-scale-with-confluent-q1-22-launch/>
- To learn how to get started with these new features, register for the upcoming Confluent Q1 ‘22 Launch webinar: <https://www.confluent.io/resources/demo/learn-the-confluent-q1-22-launch>
- See how Confluent is helping its customers transform their businesses: <https://www.confluent.io/customers/>
- Join Confluent and apply for one of its open positions: <https://www.confluent.io/careers/>

About Confluent

Confluent is pioneering a fundamentally new category of data infrastructure focused on 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 backend operations. To learn more, please visit www.confluent.io.

The preceding outlines our general product direction and is not a commitment to deliver any material, code, or functionality. The development, release, timing, and pricing of any features or functionality described may change. Customers should make their purchase decisions based upon services, features, and functions that are currently available.

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.

View source version on [businesswire.com](https://www.businesswire.com/news/home/20220119005361/en/): <https://www.businesswire.com/news/home/20220119005361/en/>

Media Contact

Taylor Jones
pr@confluent.io

Source: Confluent, Inc.