



Confluent Launches Q3 '21 Release, Making It Easier Than Ever for Any Organization to Set Data in Motion

August 17, 2021

Quarterly releases provide an easy way to learn how Confluent goes beyond Apache Kafka to deliver a platform for data in motion that is cloud-native, complete, and everywhere

Everywhere: Cluster Linking adds new data mirroring experience for near-real-time disaster recovery

Complete: New fully managed connectors for Salesforce and Microsoft Azure help deliver real-time business intelligence and richer customer experiences

Cloud-native: Infinite Storage, now available for Google Cloud customers, retains real-time and historical data without limits

MOUNTAIN VIEW, Calif.--(BUSINESS WIRE)--Aug. 17, 2021-- [Confluent, Inc.](#), the platform to set data in motion, today announced the Confluent Q3 '21 Release. This is Confluent's first quarterly release, and it features developments that help organizations reliably share data between different environments, seamlessly integrate with business-critical applications, and cost-effectively store data needed for next-generation, digital customer experiences and data-driven backend operations.

With the rise of digital native companies, everyone expects real-time, in-the-moment experiences—from the way businesses operate behind the scenes to the way people complete everyday tasks like banking and grocery shopping. But too often, the data needed to fuel applications for this new paradigm cannot keep up. Existing data architectures were designed for data at rest, so they simply aren't equipped for continuous, real-time flows of data.

"At the incredible pace of innovation happening in every industry, ubiquitous access to data in real time is now more crucial than ever," said Ganesh Srinivasan, Chief Product and Engineering Officer, Confluent. "The Confluent Q3 '21 Release helps companies harness data as it's generated across more business applications and cloud environments at global scale to build a reliable, central nervous system for data in motion."

Introducing the Confluent Q3 '21 Release

With every quarterly release, Confluent will announce new capabilities for data in motion and Apache Kafka® that reflect the investments it's making to deliver a platform that is cloud-native, complete, and available everywhere. The Confluent Q3 '21 release aims to help organizations break data silos, reliably share data between different environments, and simplify how organizations use data to react, respond, and adapt in real time.

Everywhere: Cluster Linking adds new data mirroring experience for near-real-time disaster recovery

"In order to meet new architectural requirements and reduce costs, we needed a solution for migrating data and existing workloads to a new Kafka cluster," said Zen Yui, Data Engineering Manager, Namely. "We completed this migration quickly and easily using Confluent's Cluster Linking. With perfectly mirrored topic data/metadata replication, offset preservation, and support for non-Java consumers, the migration was even more simple than we expected."

Most organizations today are global and have architectures that span across continents and multiple cloud environments. Ensuring data flows consistently across these diverse environments in a way that's secure and reliable is increasingly difficult. It typically requires additional, advanced replication tools that increase architectural complexity, are notoriously hard to manage, and often delay application development.

To help overcome this challenge, [Cluster Linking](#) is now generally available, providing a simple, fully managed solution for connecting Kafka clusters and replicating events across regions and clouds. This makes it easy to establish a single source of truth for real-time data that can be used for global data sharing, disaster recovery, and easy workload migrations. As part of this release, Confluent added additional capabilities for [near-real-time disaster recovery](#). With Cluster Linking's new data mirroring experience, a single Confluent Cloud CLI or API command lets you failover to perfectly mirrored Kafka topics on a separate region or different cloud provider altogether. This is an important capability for keeping business-critical applications running in the event of outages with a public cloud provider or regional service disruption.

Complete: ksqlDB pull queries simplify stream processing architectures, making it easier to build applications powered by data in motion

"Our customers expect instant updates on their order status and what's in stock, which makes processing inventory data in real time a must-have for our business," said Chirag Dadia, Director of Engineering, Nuuly. "ksqlDB pull queries enable us to do point-in-time lookups to harness data that is critical for real-time analytics across our inventory management system. Now, we can pinpoint exactly where each article of clothing is in the customer experience."

For organizations to get the full benefit of data in motion, they need a complete set of capabilities to build applications quickly, reliably, and securely. Confluent drastically simplified the process of building data in motion applications with [ksqlDB](#), the only database purpose-built for stream processing applications. It first launched with [push queries](#), which deliver continuous streams of updates in real time, like the location of a rideshare or a retailer's current inventory. However, many applications still need to pull static, historical information like purchase history and contact information. As a part of the Q3 '21 Release, [ksqlDB pull queries](#) are now generally available in Confluent Cloud. Together with push queries, ksqlDB enables a broad class of end-to-end stream processing workloads without the need to work across multiple systems to build streaming applications. Organizations now have a

single, easy-to-use solution for building innovative customer experiences powered by data in motion.

Complete: New fully managed connectors for Salesforce and Azure help deliver real-time business intelligence and richer customer experiences

Confluent has expanded its [library of fully managed Kafka connectors](#) that seamlessly integrate business-critical applications and data systems across an entire organization. These new connectors help modernize and future-proof data infrastructure, giving organizations the flexibility to share data broadly and build everywhere across any environment. The [Salesforce Platform Events Source](#) connector enables organizations to unlock valuable customer data and share it with downstream data warehouses and applications to improve business intelligence and analytics and enrich customer experiences. The fully managed [Azure Cosmos DB Sink](#) connector helps companies migrate to a modern, cloud-native database to enable high performance and automatic scaling, unlocking real-time use cases like IoT, retail services, and online gaming at massive scale.

Cloud-native: Infinite Storage, now available for Google Cloud customers, retains real-time and historical data without limits

To avoid the costly operational overhead of managing data in motion at massive scale, Confluent re-architected Kafka to work seamlessly in the cloud with cloud-native traits developers expect, including elasticity, high availability, and cost-effectiveness. [Infinite Storage](#) is now generally available for [Google Cloud](#), after initially launching for [AWS](#). Organizations can now retain all the real-time and historical data they need without the hassle of pre-provisioning additional infrastructure or running the risk of paying for any unused storage. With real-time, data-driven applications based on a limitless supply of both current and historical information in the same place, businesses are able to build more engaging customer experiences and augment software-driven backend operations to take faster, more informed action.

Get started now

For a deep dive into the Confluent Q3 '21 Release, [register for a four-part webinar](#) series that provides quick tutorials on how to get started with all the latest capabilities. More detail on each feature can also be found on [the Confluent blog](#).

Additional Resources

- See how Confluent is helping its customers transform their businesses: <https://www.confluent.io/customers/>
- Learn more about Confluent: <https://www.confluent.io/>
- 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/20210817005341/en/): <https://www.businesswire.com/news/home/20210817005341/en/>

Lyn Canilao
pr@confluent.io

Source: Confluent, Inc.