

CONFLUENT

# Getting Real-time Analytics and Insights Has Never Been Easier



Confluent helps you:

- **Move from batch processing to streaming ETL at scale**

Design high-performance and secure real-time data pipelines at enterprise-wide scale by building a streaming architecture across all environments—hybrid, multi-cloud, or on-prem

- **Reduce compute, storage, and operational costs**

Minimize total cost of ownership (TCO) for ETL workloads by decreasing infrastructure requirements and offloading the operational burden to a fully managed, cloud-native service

- **Deliver rich insights and real-time analytics faster**

Unlock new business insights in real-time by transforming and enriching your data while in transit and flexibly forwarding it to an analytics destination of choice

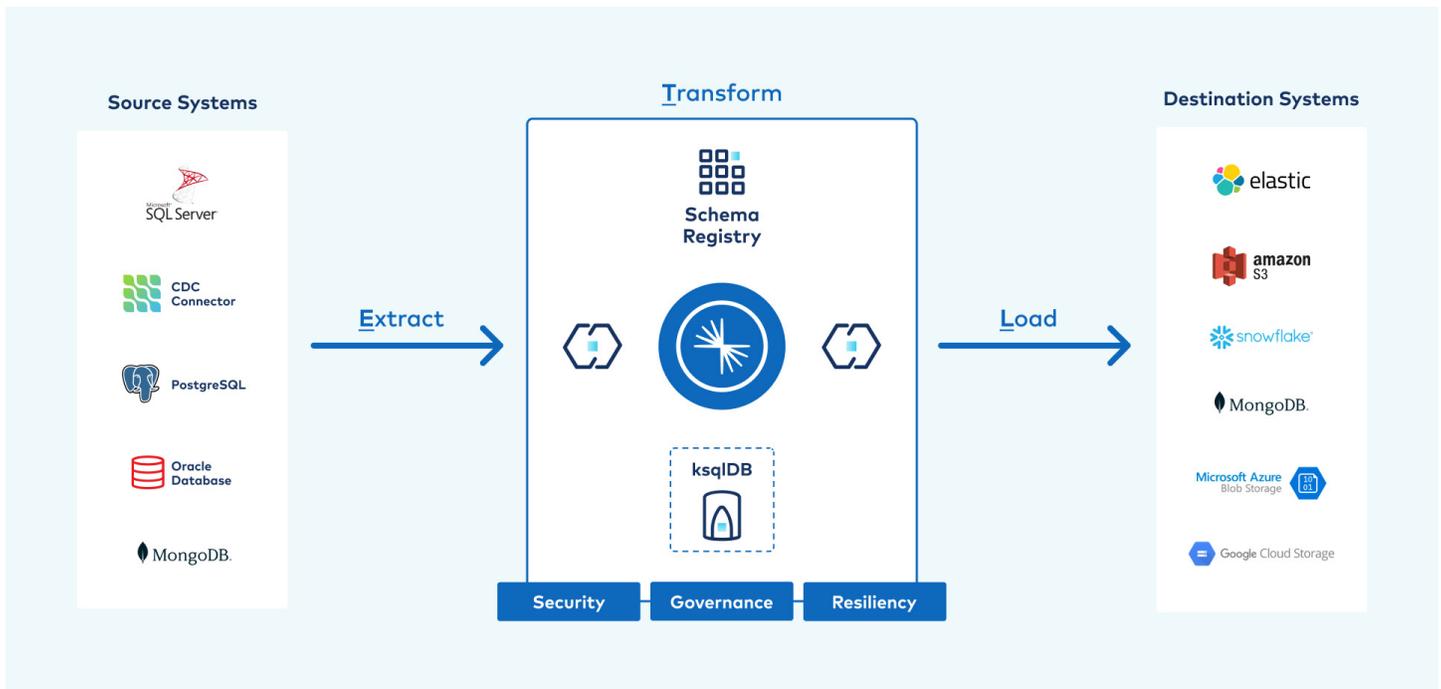
In a world where customers expect immediacy, real-time data can be the differentiator between business success and failure. It is no surprise then that organizations across the globe are prioritizing real-time initiatives as a means of gaining an edge over their competitors and improving operational efficiencies.

To enable these real-time data initiatives, data needs to not only be ingested, but processed in real time. This typically requires performing analytic computations on streaming data streams, otherwise known as streaming analytics, to extract immediate insights from newly ingested data. Yet, while the uses and demands for data have changed, organizations' underlying infrastructure have not. Most organizations still use batch-oriented extract, transform, load (ETL) tools as a core component of their data integration process. However, the need for immediacy means the large latency between gathering data and preparing it for analysis must be addressed.

This paradigm shift in data processing is why we are seeing more organizations move towards streaming ETL, which can process and take action on data as soon as it arrives in an ETL pipeline. Confluent enables simple, scalable data streaming pipelines with 120+ pre-built connectors for real-time integration between source and destination systems, in-flight processing of data streams with ksqlDB, and a robust set of security, governance, and resiliency features to meet enterprise requirements for distributed mission-critical workloads. Confluent also enables customers to go beyond just real-time integration between data systems, supporting real-time stream processing and analysis for real-time decision making and powering modern applications.

# Move from batch processing to streaming ETL at scale

Confluent offers a robust solution to enable high-performance, real-time data pipelines at scale across all of your different environments – hybrid, multi-cloud, or on-prem. By leveraging Confluent’s 120+ pre-built connectors, you can easily integrate your production systems to your chosen data warehouse or cloud storage, saving you months of custom development time. Beyond simplified integrations, Confluent also offers ksqlDB to make in-flight transformations like filtering, joining, and aggregating high-throughput data streams with lightweight, SQL-based syntax easy. For those who operate across multiple environments, enable secure and real-time movement of data across regions, clouds, and hybrid environments with Cluster Linking to create a consistent, real-time view of all data throughout your organization. Plus, with native stream governance, you can discover, understand and trust your data streams as it moves across your enterprise.



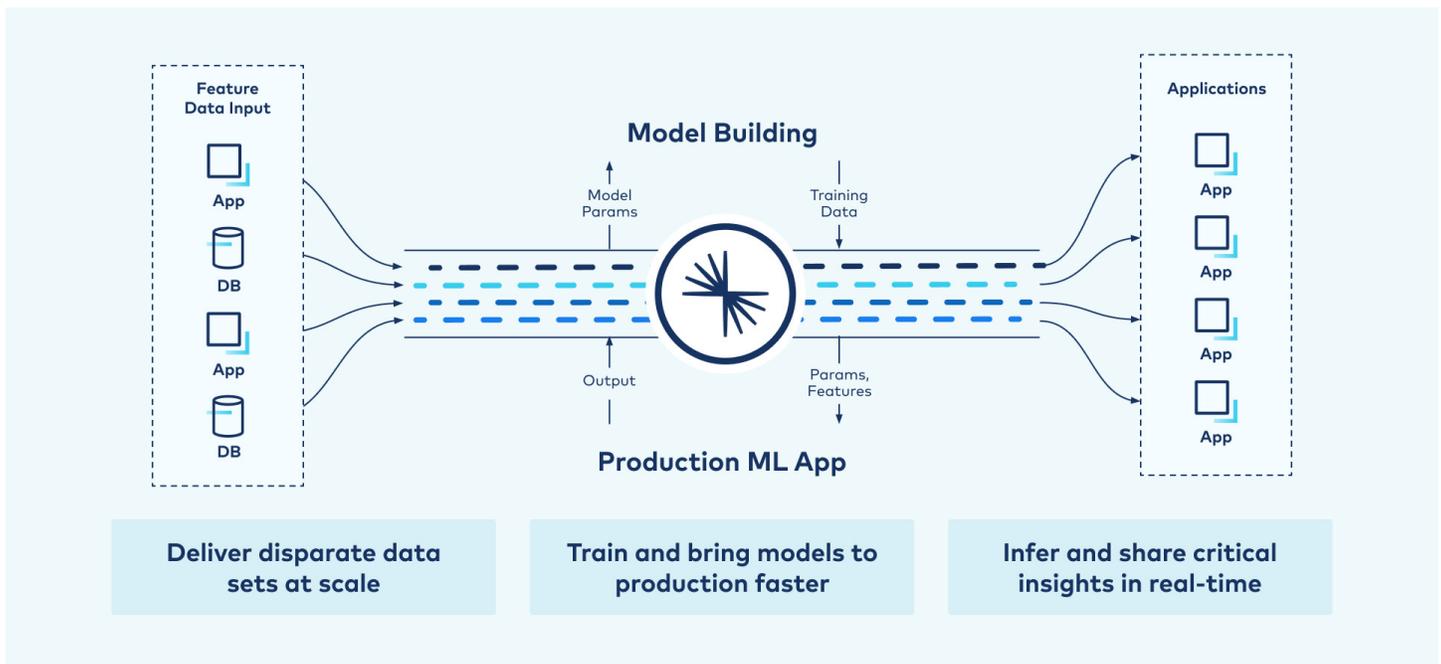
# Move from batch processing to streaming ETL at scale

Confluent minimizes total cost of ownership (TCO) for ETL workloads by decreasing infrastructure requirements and offloading the operational burden to a fully managed, cloud-native service. By using Confluent as a real-time, persistent pipeline, you can reduce the high compute and storage requirements typically associated with ETL workloads. In addition, you can eliminate the need to store and process uncleaned data in expensive downstream systems by prepping and processing the data upstream and leveraging Confluent’s infinite storage to store data in perpetuity. Using Confluent’s cloud-native service also means never needing to manually configure, provision, or deprovision any VMs typically needed to run batch ETL processes. Rather than writing custom scripts or integration tools, you and your team can leverage Confluent’s enterprise-ready out-of-the-box functionalities making it easy to go from idea to proof of concept. Confluent also makes it easy to scale elastically to meet any workload and completely abstracts away tedious activities like patching and upgrades.

# Deliver rich insights and real-time analytics faster

Confluent delivers rich insights and real-time analytics faster by processing a highly diverse set of data streams instantaneously. Confluent provides sophisticated stream processing capabilities to curate, enrich, transform and normalize data in real time, all through simply SQL syntax. Then teams can apply simple business logic such as stateless filtering or stateful aggregations, or design complex business rules with user-defined functions.

Confluent also enables forward-compatibility and easily processes a highly diverse dataset by using ksqlDB to transform data at ingestion, avoiding manual errors and poor data integrity downstream. So now you and your team can focus on delivering value to the business by building ML/AI models that detect anomalies or make predictions on data in flight (e.g., fraud detection) or triggering immediate action on business-impacting events as they occur (e.g., notifications for flight delays). In addition to streaming ETL, ksqlDB can be used to create materialized views to serve real-time dashboards and reporting (e.g., inventory management) without ever needing to move the data to another system.



## 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. To learn more, please visit [www.confluent.io](http://www.confluent.io).

To learn about how Confluent can help you modernize your applications or build new streaming applications, check out our [streaming ETL solutions page](#).