

Discover, Understand, and Trust Your Data Streams on Confluent Cloud

Safely expand the value of streaming data pipelines throughout your entire organization with the only governance suite purpose-built for the intricacies of data in motion. **Stream Governance** establishes trust in the real-time data moving throughout your business and delivers an easy, self-service experience for more teams to discover, understand, and put these streams to work. Go beyond expertise-silos and accelerate data streaming initiatives to establish a central nervous system of real-time data within your business without bypassing controls for risk management or regulatory compliance.



Deliver trusted, high-quality data streams to the business



Increase collaboration and productivity with self-service data discovery



Understand complex data relationships and uncover more insights

Why Stream Governance?

With the sharp rise of real-time data, Apache Kafka®, and event-driven systems, the need for organizational governance over data streams is growing fast. While tools for building and maintaining long-term, compatible streaming pipelines like Schema Registry have been available for some time, the ability to safely and easily share these streams across teams for more widespread use has not. This need is felt most by businesses deploying distributed, event-driven microservices built by small, disparate teams of data streaming experts.

Stream Governance provides businesses with a modern governance solution, built for data streaming, that allows them to operate at scale around a central nervous system of real-time data. Individual users across teams have access to the data they need through self-service discovery in a system with the controls required to ensure long-term data compatibility, risk management, and regulatory compliance. When distributed teams can explore, understand, and trust data streams together they can harness its full value. Businesses now have the key to fostering the collaboration and knowledge sharing necessary to become an event-centric business while remaining compliant within an ever-evolving landscape of data regulations.

Feature Pillars



Stream quality

Set and enforce the global rules by which the entire system operates in order to control what data gets in and what data does not—all in the spirit of maintaining high data integrity



Stream catalog

Unlock the full value within your data with self-service discovery tools that allow teams to classify, organize, share, and find the high-quality data streams they need



Stream lineage

Comprehend and troubleshoot complex data streams easily with interactive, end-to-end lineage maps provided with single-click export for regulatory compliance

Stream quality: Deliver trusted, high-quality data streams to the business

Confluent's stream quality tools, provided across 28 global regions with a 99.95% uptime SLA available, enable teams to deliver a scalable supply of trusted data streams throughout the business, enabling reliable delivery of mission-critical applications, confident decision making, and a simplified design for data standards.

Schema Registry allows teams to define and enforce universal data standards that enable scalable data compatibility while reducing operational complexity. Schemas shared across cloud and hybrid environments sync in real time with **Schema Linking**. Avro, JSON and Protobuf serialization formats are supported.

Schema Validation, enabled at the topic-level, ensures broker/registry coordination by verifying that schemas tied to incoming messages are both valid and assigned to the specific destination topic in order to publish.

Standardizing around well-defined and agreed upon schema structures allows teams to develop resilient pipelines prepared for safe, compatible evolution over time. And with the Schema Management UI, users can easily learn about all schemas that exist, how they are defined, and where they are used.

Stream catalog: Increase collaboration and productivity with self-service data discovery

Available through the Confluent Cloud UI, REST API, and GraphQL API, **Stream Catalog** allows users across teams to collaborate within a centralized, organized library designed for sharing, finding, and understanding the data needed to drive their projects forward quickly and efficiently.

Built and organized through automatic metadata collection, this catalog includes topics, schemas, fields in schemas, clusters, connectors, and more. It's like a digital library for data streaming assets allowing any user—experienced with Kafka or not—to search for what they need, find what's already been built, and put it to use right away.

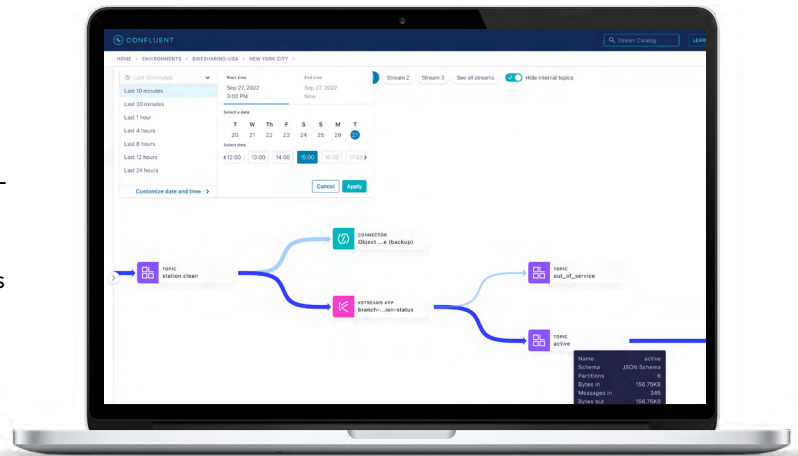
With custom classifications and tagging, business metadata options, and open-form entity descriptions, teams and individual users alike can constantly increase the value of the catalog by adding contextual, business-specific details to data and improve data discovery throughout the entire business.

Stream lineage: Understand complex data relationships and uncover more insights

Moving forward with updates to mission-critical applications, troubleshooting production workloads, answering questions on important subjects like data regulation & compliance, or kicking off a new project with sensitive data requires an easy means of comprehending the big picture journey of streaming data.

Stream Lineage provides a GUI of data streams and data relationships with both a bird's eye view and drill-down magnification for answering questions like: Where did data come from? Where is it going? Where, when, and how was it transformed? When and how was a pipeline last updated? Live metrics and metadata inspection embedded directly within lineage graphs enables teams to learn quickly and make more informed decisions. Historical, point-in-time insights and search across graphs simplify maintenance of mission-critical data streams.

And with full insight into where and how their data is moving, including single-click export of end-to-end streams, businesses are able to ease data reporting for audit purposes and increase visibility during procedures like workload migrations.



Getting Started

Stream Governance is provided fully managed within Confluent's cloud-native data streaming platform, **Confluent Cloud**. New signups receive \$400 to spend during their first 30 days—no credit card required.



"Running a modern bank serving fast growing enterprises requires data in motion we can trust. Schema Registry allows us to provide customers with reliable, real-time experiences for tasks like loan processing while we maintain focus on furthering innovations. We're excited about new discovery capabilities within Stream Governance which will enable us to expand data in motion across teams and operate like a truly event-centric business."

— ANDREAS PIEFKE, HEAD OF CLOUD ARCHITECTURE, JUDO BANK