

CONFLUENT PLATFORM

Rich pre-built ecosystem

[Confluent Platform](#) offers a rich pre-built ecosystem of over 100 Apache Kafka® connectors and a schema registry to rapidly and reliably build event streaming applications around Kafka.



Instantly connect popular data sources and sinks



Enable application development compatibility



Deploy confidently in production

Why a rich pre-built ecosystem?

Kafka developers should spend the majority of their time building an ecosystem of event streaming applications that will change their business, and less time figuring out the inner workings of Kafka. To build an event streaming ecosystem you need to leverage data that resides in existing systems, such as databases, data lakes or mainframes. However, developing your own connectors is time-consuming, and using open source connectors can be risky for enterprise users. Additionally, enabling compatibility across event streaming applications usually requires lengthy and complex coordination among developers, which can ultimately be error-prone and risky.

Confluent Platform enables fast, broad and reliable Apache Kafka connectivity through pre-built Kafka connectors and programmatic schema management.

Features



Pre-built Kafka Connectors

Confluent develops and works with partners who develop enterprise-ready connectors based on the Kafka Connect framework. Connectors are supported by either Confluent or our partners. A portion of them are available as managed connectors with [Confluent Cloud](#).



Confluent Hub

[Confluent Hub](#) is an online marketplace to easily browse, search, and filter for connectors and other plugins that best fit your data movement needs for Kafka.



MQTT Proxy

MQTT Proxy delivers Kafka-native connectivity into IoT devices without the need for intermediate MQTT brokers, eliminating the additional cost and lag. MQTT Proxy accesses, combines and guarantees that IoT data flows into Kafka without adding additional layers of complexity.



Schema Registry

Schema Registry is a central repository with a RESTful interface for developers to define standard schemas and register applications to enable compatibility. Schema Registry is available as a software component of Confluent Platform or as a managed component of [Confluent Cloud](#).



“The addition of interfaces missing from the core open source Kafka such as the REST Proxy, Replicator, Control Center and the long list of supported connectors. Confluent really makes their Kafka offering very complete”

Architect | [Gartner Peer Insights](#)

Solution

Instantly connect popular data sources and sinks

100+ pre-built Confluent enterprise connectors

The catalog of fully supported connectors includes JDBC, HDFS, AWS S3, Elasticsearch, MongoDB, Salesforce, Debezium, MQTT, and many more. Some connectors are also available as managed components of [Confluent Cloud](#), such as AWS S3, Google GCS & BigQuery, Azure Blob.

Easily find Kafka Connectors with Confluent Hub

To simplify how you leverage the Kafka Connect connector ecosystem, we offer [Confluent Hub](#), an online marketplace to easily browse, search and filter connectors to find the one that fits your needs.

Connect MQTT IoT data sources

Enable IoT data to flow into Kafka without adding additional layers of complexity. MQTT Proxy delivers Kafka-native connectivity into IoT devices without the need for intermediate MQTT brokers, thereby expanding the event streaming platform into new enterprise data sources and business applications.



Enable application development compatibility

Develop using standard schemas

Store and share a versioned history of all standard schemas, and validate data compatibility at the client level. Schema Registry supports Avro, JSON and Protobuf serialization formats.

Reduce operational complexity

Schema Registry reduces operational complexity in the application development cycle, because it eliminates the need for complex coordination among developers. Need to add a new column to a downstream database? You don't need an involved change process and at least 5 meetings to coordinate 20 teams.

Deploy confidently in production

Scale Kafka Schemas reliably

Schema Validation delivers a programmatic way of validating and enforcing Schema Registry schemas directly on the Kafka broker and with topic-level granularity. It provides greater control over data quality, which increases the reliability of the entire Kafka ecosystem.

Manage the Kafka ecosystem centrally

Simplify management for production environments using Control Center as the GUI:

- Manage multiple connectors: add, edit and delete connectors across multiple Connect clusters
- Manage every schema: create, edit and view topic schemas, compare versions, and enable Schema Validation when creating new topics

Confluent Platform. Enterprise event streaming platform built by the original creators of Apache Kafka. For more information, please visit [confluent.io](#). To contact us, visit [confluent.io/contact](#). For detailed product specifications, please refer to our [documentation](#).