CONFLUENT PLATFORM

Dynamic performance and elasticity

Confluent Platform enables unparalleled Apache Kafka® performance and elasticity through automated partition rebalancing, infinite data retention, and the decoupling of the compute and storage layers.

Optimize throughput across brokers dynamically

Retain infinite data on Kafka

Elastically scale Kafka clusters

Why dynamic performance and elasticity?

Kafka was designed to be massively scalable and provide high throughput for significant volumes of messages. However, Kafka operators still face several common challenges when optimizing Kafka's performance and scaling their cluster as event streaming spreads across their organizations, including diminished throughput, increased storage costs, and a lengthier and more inefficient scaling process when adding brokers to the cluster.

Confluent Platform ensures that your Kafka cluster stays performant and can quickly scale to meet the needs of your organization. By optimizing the overall performance for allocated infrastructure, Confluent Platform minimizes infrastructure costs for Kafka and maximizes your ROI on event streaming.

Features

Auto Data Balancer

Auto Data Balancer (ADB) provides an automated solution for partition rebalancing to optimize the cluster's throughput. ADB determines the best cluster-wide partition rebalancing plan without manual, time-consuming calculations or risk of human error. ADB can be used on an ongoing basis to optimize partition placement, or whenever brokers need to be added or removed.

Tiered Storage (preview)

Tiered Storage allows Kafka to recognize two tiers of storage: local disks and cost-efficient object stores (Amazon S3). Brokers can offload older topic data to object storage, enabling virtually infinite retention. Kafka consumers require no unique configuration to maintain transparent and performant access to messages and topics.

EURONEXT

"We have been very satisfied with Confluent Platform as the backbone of our persistence engine. The platform has been super reliable. We have stringent requirements for real-time performance and reliability, and we have confirmed — from proof-of-concept to deployment of a cutting-edge production trading platform — that we made the right decision."

Alain Courbebaisse | Infrastructure and Operations Director, Euronext

Solution

Optimize throughput across brokers dynamically

Automated partition rebalancing

Ensure your cluster's topic partitions are evenly distributed among its brokers with ADB. ADB provides fully-automated partition rebalancing to ensure your cluster's throughput is optimized without the need for time-consuming calculations or risk of manual error.

Cluster-wide rebalancing

With ADB, you can also efficiently rebalance your entire cluster rather than needing to explicitly reference each individual topic one-by-one as you would with Apache Kafka tooling.

Bandwidth throttling

Limit the impact of a partition rebalance on your cluster's throughput through ADB's ability to optimize partition movement during the rebalance and throttle its allocated bandwidth, providing your cluster with sufficient resources to continue serving producers and consumers.

Tiered Storage enables infinite data retention on Kafka.

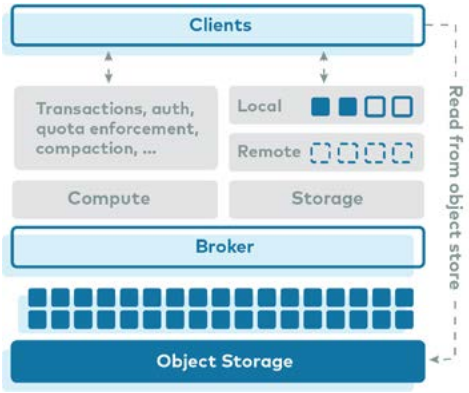

Retain infinite data on Kafka

Cost-effective object storage

Achieve dramatically longer periods of data retention for your cluster without the significant increase in operating costs. By leveraging Tiered Storage, you can back up data for replay in the future or use it to make Kafka the central nervous system of your organization.

Built to meet regulatory requirements

Tiered Storage also allows you to achieve the regulatory compliance for data retention and durability requirements specific to your industry, without needing to build additional infrastructure into your architecture.

Elastically scale Kafka clusters

Decoupled compute and storage

Scale storage resources without having to scale compute resources, and vice versa, by effectively decoupling compute and storage resources allocated to your cluster through Tiered Storage.

Immediate broker scaling

Minimize the time needed to rebalance partitions when adding new brokers. Because less data is stored on the broker when using Tiered Storage, smaller partitions need to be moved to newly added brokers using ADB, reducing rebalancing times from potentially hours down to a few seconds.

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**.