

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 data streaming.

Features

Self-Balancing Clusters

Self-Balancing Clusters automate partition rebalances to optimize Kafka’s throughput, accelerate broker scaling, and reduce the operational burden of managing a large cluster. Partition rebalances are completed quickly and without any risk of human error.

Tiered Storage

Tiered Storage allows Kafka to recognize two tiers of storage: local disks and cost-efficient object stores. 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.



"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 COURBEBASSE, INFRASTRUCTURE AND OPERATIONS DIRECTOR, EURONEXT

Solution

Optimize throughput across brokers dynamically

Automated partition rebalancing

Ensure your cluster's workload is evenly distributed among its brokers with Self-Balancing Clusters. Self-Balancing Clusters monitor the skew of resource utilization across your brokers and continuously reassign partitions to optimize your cluster's performance and balance.

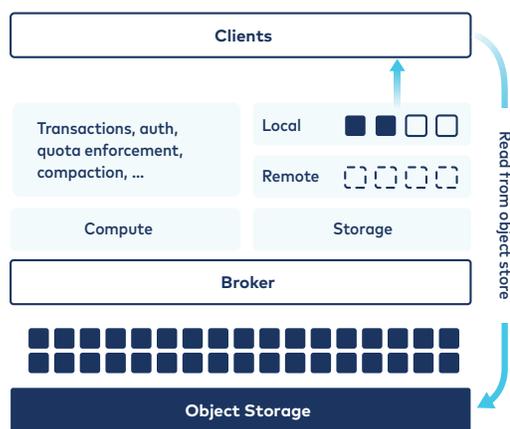
Accelerated broker scaling

With Self-Balancing Clusters, you can add or decommission brokers and a partition reassignment will automatically be triggered, providing you with even greater cluster elasticity.

Reduced operational burden

Eliminate the need for complex math and the risk of human error that partition reassignments typically entail. By leveraging Self-Balancing Clusters, Kafka operators no longer have to spend their time manually reshuffling massive loads of data across their cluster.

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 performance benefits

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 Self-Balancing Clusters, reducing rebalancing times from potentially hours down to a few seconds.

Confluent Platform. Enterprise platform for data in motion, 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](#).