

CONFLUENT PLATFORM

# Global resilience

[Confluent Platform](#) offers a comprehensive solution for building Apache Kafka® clusters that are disaster-resilient, highly available, and easily migratable across multi-cloud environments.



**Minimize disaster recovery complexity**



**Achieve faster recovery time**



**Migrate data across multi-cloud environments**

## Why global resilience?

Kafka users have the same high expectations for durability and availability that they do for all data stores. To leverage Kafka in production use cases, you need to ensure crucial data is not lost, downtime is minimized, and compliance and legal requirements are met. Achieving these objectives requires your Kafka deployments to stretch across multiple data centers around the globe, which can create significant operational complexity and costs. Enterprise users also have requirements around connecting on-prem clusters to cloud environments, migrating data between cloud providers, and creating hub-and-spoke architectures.

**Confluent Platform** offers capabilities to stretch Kafka clusters and migrate Kafka data across geographies, enabling you to make the event streaming platform the central nervous system of your business at a global level.

## Features



### Multi-Region Clusters

For disaster recovery automation, Multi-Region Clusters allow you to run a single Kafka cluster across multiple data centers, providing your enterprise with the operational simplicity of stretch cluster architecture without the typical tradeoff between throughput and disaster resilience.



### Replicator

For data migration use cases, such as aggregating clusters to a centralized location, migrating clusters across regions, or building hybrid cloud environments, Replicator provides asynchronous replication built on the Kafka Connect framework.



“We love open source, but at the same time we’re not a startup. We’re a large financial institution that works with world-class organizations, and we need services that make it easier for us to sleep at night. Confluent Platform is very reliable; it’s never down. It has become our backbone.”

Kaspar Situmorang | Executive Vice President, Bank Rakyat Indonesia

## Solution

### Minimize disaster recovery complexity

#### Simplify deployment

With Multi-Region Clusters, you can leverage Kafka’s internal replication engine to stretch one Kafka cluster across multiple data centers in different regions, ensuring your data is backed up without needing to deploy a separate Connect-based replication tool like MirrorMaker 2.

#### Automated client failover

Multi-Region Clusters also enables automated client failover, so your producers and consumers instantly failover to the recovery site upon failure without your Kafka operators needing to coordinate across your multiple developer teams or handle DNS reconfigurations.

### Achieve faster recovery time

#### Minimize recovery time objectives

Build your mission-critical applications with high uptime SLAs on top of Kafka rather than legacy, batch-oriented systems. With support for automated client failover, Multi-Region Clusters enable Recovery Time Objectives (RTO) measured in as little as a few seconds rather than hours.

#### Eliminate data loss for high priority topics

Ensure Recovery Point Objectives (RPO) of 0 for topics with your most sensitive data. Multi-Region Clusters allow you to choose between synchronous or asynchronous partition replication on a per topic basis.

### Migrate data across multi-cloud environments

#### Ensure mission-critical reliability

Leverage the reliability of Replicator for your replication use cases, such as:

- Aggregating several clusters together into one
- Migrating entire clusters to a preferred environment
- Bridging clusters across two different cloud providers
- Replicating clusters across different continents

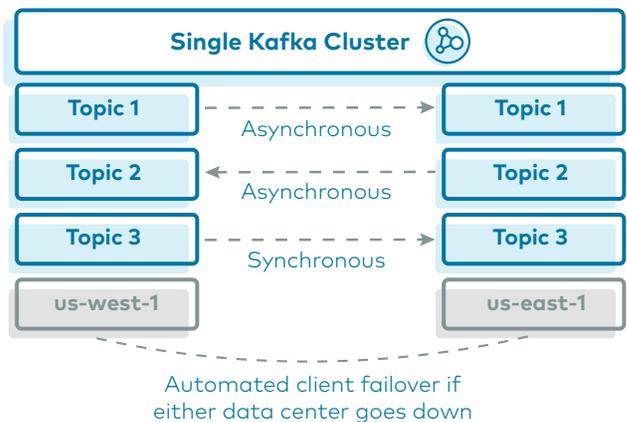
#### Bridge to Confluent Cloud

Bridge your self-managed clusters to a fully-managed service in Confluent Cloud using Replicator, providing your organization a path to offloading cluster management to the world’s foremost Kafka experts.

#### Control Center integration

Replicator comes integrated into the Control Center, so you can easily monitor the replication process and identify any potential issues prior to a disaster event.

Multi-Region Clusters run a single Kafka cluster across multiple data centers.



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