WHY CONFLUENT VS. MSK

Future-proof your investment with a fully managed, cloud-native, and complete data streaming platform

Built by the original creators of Apache Kafka®

Data in Motion is a key requirement for the modern data stack that enables organizations to connect, process, and react to their data streams to power applications and operations across their business in real time. Founded by the original creators of Apache Kafka®, Confluent has harnessed the power of the cloud to build a data streaming platform that's 10x better than Apache Kafka, helping customers reach 10x the elasticity, storage, and resiliency benefits vs. Open Source Kafka.

AWS has taken Open Source Kafka and hosted it in their cloud to offer two different products: AWS MSK and MSK Serverless.

AWS MSK (Provisioned cluster, GA at May 2019)

A hosted offering for Kafka which only offloads the infrastructure layer, requiring users to manually manage and scale cluster capacity and patches/upgrades.

MSK Serverless (Serverless cluster, GA at May 2022)

A serverless offering that manages and scales cluster capacity. However, MSK Serverless has low limits on partitions/retention, making it difficult to support production-scale use cases, while lacking many mission-critical components, such as uptime SLAs, connectors, stream processing, Kafka-specific security & governance, and more.

SecurityScorecard

"Since we built Horus, our global IPv4 scanning platform on top of Confluent, we've saved over a million dollars compared to open source Kafka or MSK. Business resilience and ensuring no disruption to delivering customer value, all of that is enabled by having a system like Confluent that works securely and reliably to do data streaming."

Jared Smith | Senior Director, Threat Intelligence, SecurityScorecard

About Confluent

As the foundational platform for data in motion, Confluent's cloud-native offering is designed to be the intelligent connective tissue enabling real-time data from multiple sources to constantly stream across the organization.

As an industry leader, Confluent has been recognized with a Top 10 Ranking on Forbes Cloud 100, Google Cloud Technology Partner of the Year Award, JPMorgan Chase Hall of Innovation, and Bank of America Recognition for Enterprise Technology Innovation.

Why is this a problem?

1. Customers need to choose between two cluster types, each with its own set of drawbacks.

2. Lack of out-of-the-box tools to migrate data/applications from one cluster type to the other – not to mention the operational overhead this would entail.

3. If customers have non-AWS environments on other clouds or on-premises, they're on their own.

That's why so many customers turn to Confluent to help eliminate their operational burdens and risks with a complete, enterprise-grade data streaming platform to power a central data backbone in real-time across their entire business.

What makes Confluent different?

Only Confluent provides a truly cloud-native and complete Kafka service ready to deploy, operate, and scale in a matter of minutes, while offloading all operational complexity, burden, and risk to the world's foremost Kafka experts to save up to 60% on your data streaming TCO.

Eliminate your ops burden and modernize your data architecture with a truly **cloud-native** Kafka solution

Future-proof your investment with a fully managed, cloud-native solution that will support scaling mission critical use cases across your org instead of settling for a short-term solution that is only fit for low scales. MSK still requires you to bear manual tasks like capacity planning / sizing, Kafka patching / upgrades, scaling, and more.

Confluent has completely re-architected Kafka from the ground-up, to provide teams with a truly cloud-native experience that delivers an elastically scalable and globally available product that is ready to deploy, operate, and scale in a matter of minutes.

Confluent's solution is cloud-native:

- **Elastic scalability** that automatically scales up and down to meet demand
- **High availability** with our 99.99% uptime SLA that covers both the software and underlying infra
- **Infinite storage** to cost-effectively store data without growing compute

Accelerate time-to-value and reduce total cost of ownership with a **complete** platform for all of your data-in-motion needs

Confluent provides a holistic set of integrated, enterprise-grade capabilities, including connectors, stream processing, security & data governance, and disaster recovery – so your teams can focus on efforts that differentiate your business instead of low-level infrastructure. Best of all, Confluent's solution helps customers save up to 60% in TCO.

MSK lacks pre-built connectors to integrate Kafka with your other data systems – users need to bring and self-manage their own connectors. Custom-built connectors require maintenance, and Kafka community connectors are not covered by AWS technical support. Users also need to self-provision and manage stream processing tools in order to enrich any data in-flight. Finally, MSK lacks Kafka-specific security & governance, such as resource-level RBAC and data quality/lineage/catalog tools purpose-built for Kafka.

Confluent's solution is complete:

- 120+ battle-tested, **pre-built connectors**
- **Enterprise-grade security** & **governance**
- **ksqlDB** for easy stream processing
- **Stream Designer** rapidly build pipelines end-to-end with an easy to use, visual designer

Seamlessly connect your data and applications **everywhere** they reside across hybrid and multi cloud architectures

Confluent exists everywhere your applications and data reside, allowing you to leverage a fully managed service on all leading public clouds and self-managed software for your private infrastructure. Best of all, you can seamlessly connect it all together in real time with Cluster Linking to create a consistent data fabric across your business.

Amazon MSK can be hosted on AWS only, leaving you on your own for any hybrid or multi-cloud deployments.

Confluent's solution is everywhere:

- **Fully managed cloud service** on all leading public clouds
- **Self-managed software** for on-premises workloads
- **Cluster Linking** to seamlessly link all your clusters together across any environment

RevLifter® "We want our engineers to focus on making the RevLifter platform world class, not managing infrastructure and worrying about Kafka outages. With our previous provider, we still had to consider operational aspects such as brokers and storage – with Confluent's platform, we get true elasticity which is critical for scaling on big retail events such as Black Friday or when signing on large new clients – all backed by a higher SLA."

Ian Compton | Technology Director, RevLifter

Our Customers

citi BOSCH Square DB instacart Expedia KeyBank

Walmart CapitalOne nuuly Domino's RBC The Home Depot ticketmaster Humana

Features Comparison

	Confluent	MSK	MSK Serverless
Core Kafka			
Apache Kafka Latest stable release	●		
Storage Compaction A stream (topic in Kafka) is scanned periodically to remove any old events that have been superseded by newer events that have the same key	●	◐	◐
Exactly-once Guarantees If a producer retries sending a message, the message is still delivered exactly once to the end consumer	●	◐	◐
3x Replication Factor	●	◐	◐
Cloud Native / Fully Managed			
Instant Provisioning Cluster provisioning and configuration	●	◐	◐
Elastic Scalability Scale up and down from 0 to GBps without over-provisioning infra	●		
Infinite Storage Cost-effectively retain data at any scale without growing compute	●		
No-touch Upgrades Always the latest stable version with non-disruptive rolling upgrades	●		N/A
Complete Platform			
Connect			
Managed Connect Infrastructure	●		◐
Fully managed connectors 120+ pre-built connectors, including S3, Elastic, MongoDB, Snowflake, Lambda, SFDC, and more	●		
Stream Processing			
ksqlDB/Stream Processing Real-time stream processing and materialized views with lightweight SQL syntax	●		
Lambda Event-driven compute service that lets you run code for virtually any app or backend service without provisioning or managing servers	●	◐	◐
KDA Kinesis Data Analytics	●	◐	◐
Stream Designer Rapidly build pipelines end-to-end with a easy to use, visual builder that's extensible with SQL	●		
Data Governance			
Schema Registry Central registry to ensure data compatibility	●	◐	◐
Schema Validation Broker-side Schema Validation enables the broker to verify that data produced to a Kafka topic is using a valid schema ID in Schema Registry	●		
Stream Catalog Self-service data discovery to search, classify, & organize your data streams	●		
Stream Linage Understand data lineage with interactive, end-to-end mapping of your data streams	●		
Security			
Encryption (at rest and in transit)	●	◐	◐
Private Networking	●	◐	◐
RBAC Role Based Access Control with resource-level granularity (e.g., Kafka topics, connectors, schemas)	●	◐	◐
OAuth OAuth is an open-standard protocol that grants access to supported clients using a temporary access token	●		
Audit Logs Track user and application activities to detect security threats & anomalies and ensure compliance	●	◐	◐
SASL/[PLAIN or SCRAM] Salted Challenge Response Authentication Mechanism addresses the security concerns with traditional mechanisms that perform username/password authentication	●	◐	
Built-in Compliance SOC 1/2/3, ISO 27001, GDPR/CCPA/HIPAA Readiness, PCI DSS, CSA Star Level 1, Financial Services Regulation Compliance	●	◐	◐
Everywhere / Global Availability			
AWS	●	◐	◐
Microsoft Azure	●		
Google Cloud	●		
On-premises / Private Cloud	●		
Cluster Linking Seamlessly share topics across clusters in any environment to support hybrid/multi-cloud architectures	●		
Support			
SLAs covering both Kafka software and infra	●		

Get Started with Confluent Sign up for Confluent and start streaming in minutes at confluent.io/get-started.

● ◐ Comparable functionality available