

How to Build Real-Time Products, Faster

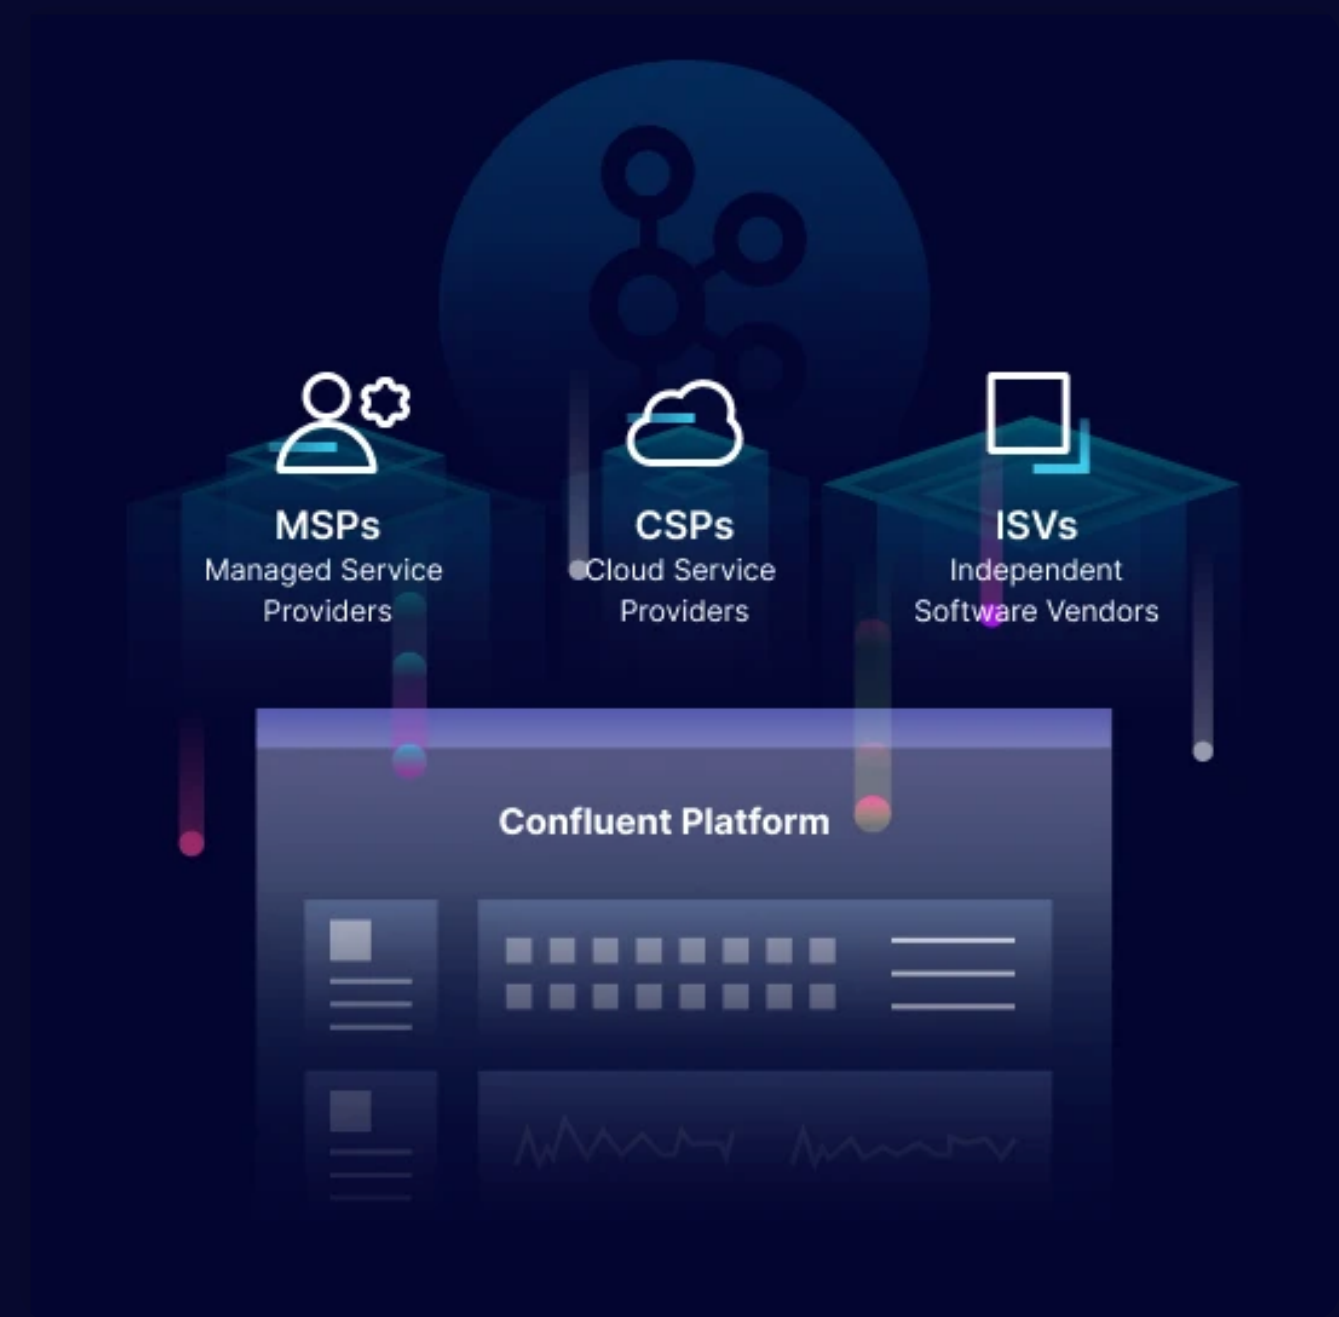
Why Businesses Choose Confluent's OEM Program Over DIY Apache Kafka®

Contents

- 6 Executive Summary
- 8 Introduction: Developing Products in the Age of Real-Time Everything
- 12 Understanding the True Cost of Apache Kafka®
- 16 Open Source vs Enterprise-Grade Data Streaming
- 20 A Deep Dive: Confluent Platform & WarpStream
- 26 The OEM Program: Grow Faster With Confluent
- 28 Start Building Real-Time Products With Enterprise-Grade Data Streaming

Confluent’s OEM Program

Grow Faster With Enterprise-Grade Data Streaming



Executive Summary

IN TODAY’S HIGHLY COMPETITIVE LANDSCAPE, real-time responsiveness is what sets businesses apart. Modern businesses need data and software services that operate on always fresh and deliver up-to-the-moment insights—anything less adds unnecessary risk.

Whatever products you build next, you need real-time data streaming and processing capabilities in your tool box. And over the last decade, more than 150,000 companies have decided on Apache Kafka® as the open standard for data streaming. Its widespread adoption is a testament to its power, flexibility, and impact in transforming how businesses operate with their data. For most, Kafka has helped define what “real time” means in the modern software era.

However, while immensely impactful, Kafka’s distributed architecture makes it inherently complex and difficult to operate. But leveraged correctly, it is hands down the ideal choice for processing large volumes of data with low latency and high fault-tolerance. That’s exactly what you need when developing for real-time payments, ecommerce personalization, patient monitoring, or any other service meeting the expectations of today’s world.

Your customers are looking to you for trustworthy solutions that help them move faster, make smarter decisions, and operate more efficiently. You can make these outcomes a reality for them with data streaming, and we can show you what it takes.

At Confluent, we know that the path to production-ready, enterprise-grade data streaming is not easy when you try to stand it up yourself.

In this ebook, we’ll show you the fast-track for launching real-time products: Confluent’s OEM Program.

As an OEM partner, you’ll bring data streaming to your product quickly and confidently with licensing to embed our complete platform with unified Apache Kafka and Apache Flink®, built and supported by the original creators of open source Kafka—the world’s foremost experts in data streaming.

Ready to learn what it takes to build real-time products with enterprise-grade data streaming?

LET’S GET STARTED.

[SKIP AHEAD TO LEARN ABOUT CONFLUENT’S OEM PROGRAM >](#)

Confluent Powers Data Streaming Across Industries



By leveraging Apache Kafka from Confluent, we ensure our RTP UPI payment platform can handle the increasing demand for real-time transactions with unparalleled efficiency and reliability. Through this partnership, we reaffirm our dedication to delivering innovative, scalable, and resilient payment solutions to clients globally, highlighting our commitment to providing cutting-edge, robust solutions worldwide.

George Sam, Co-founder and Business Head, Mindgate Solutions Pvt. Ltd.

Introduction

Developing Products in the Age of Real-Time Everything

BEING DATA-DRIVEN IS NO LONGER ENOUGH—today’s businesses have to be ready to operate and react in real time. Seconds of latency mean that frontend experiences and backend operations are based on a view of data that doesn’t reflect the current reality.

The growing demand for real-time data has shaped digital transformation and innovation over the past two decades. In pursuit of speed and agility, thousands of companies have repeatedly rebuilt fragile data pipelines, adopted messaging services, and implemented micro-batch processing and low-latency databases—each step an attempt to keep up with the pace of modern business.

DID YOU KNOW

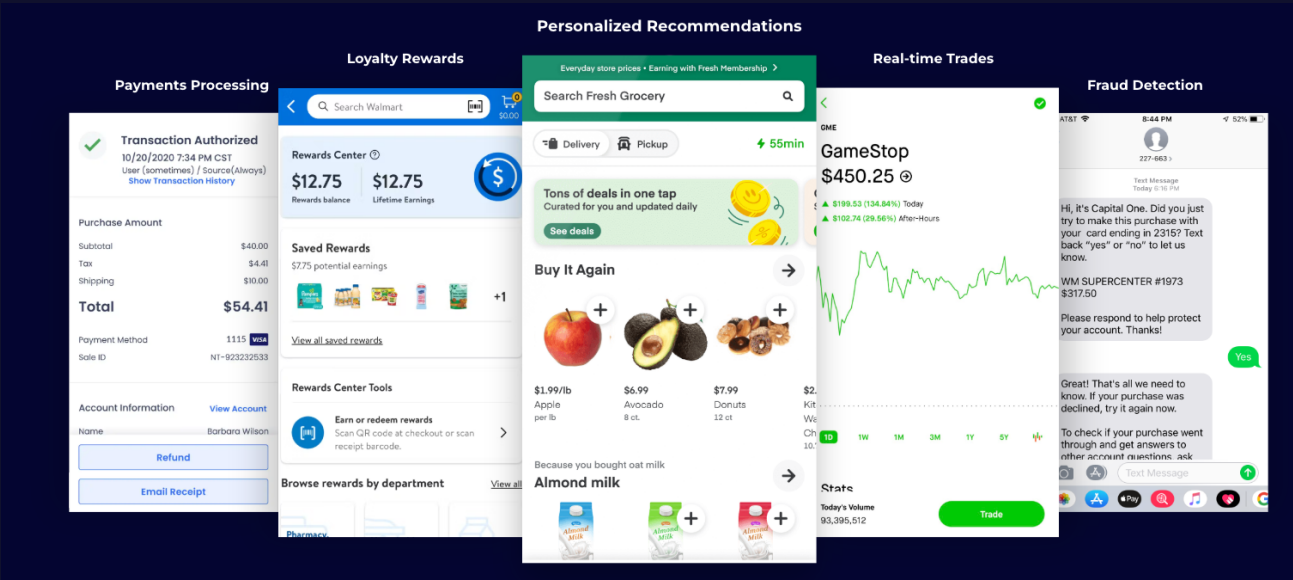
Kafka—initially created at LinkedIn to implement real-time user behavior tracking—is now the cross-industry standard for data streaming.

Today, there are more than:

- **80%** of the Fortune 100 estimated to be using Kafka
- **150,000** organizations using Kafka
- **41,000** Kafka practitioners attending 200+ global meetup groups

Even with all this dedicated expertise, Kafka often requires 2+ years to reach production at scale.

Real-Time Data Fuels Modern Business



Successfully deploying GenAI requires contextualized, trustworthy, real-time data

Whether you’re building a payments platform, ecommerce personalization features, or patient monitoring applications, you need data streaming and processing capabilities to make “real time” a reality for your business.

Today, more than 150,000 companies have adopted Apache Kafka®, an open source distributed system that has become the standard for data streaming.

Kafka started as a way for LinkedIn to ingest and apply data in real time. Once open-sourced, it quickly became the foundation for use cases disrupting every industry. Today, practitioners and technical leaders have brought over [1,000 Kafka-powered solutions](#) to life, enabling event-driven architectures, continuous real-time data processing, and instant reactions to real-world events.

Data streaming powers essential capabilities for analytics products and operational solutions today, especially in fast-moving industries like financial services, automotive & manufacturing, retail and ecommerce, and telecommunications.

Traditional, [batch-based systems](#) process stored data at periodic intervals. Adopting Kafka allows you to build software that processes continuous streams of data and application architecture capable of reacting to events as they occur.

Done right, Kafka helps you deliver the modern application architectures and real-time time solutions your customers expect like:

Financial fraud detection and risk management

Kafka allows real-time streaming and correlation of events from multiple data sources (credit card transactions, user behavior, geolocation), allowing businesses to build software that continuously monitors financial transactions for suspicious activities.

eCommerce and digital marketing personalization

Kafka streams real-time user interactions (e.g., clicks, searches, purchases). Businesses can then develop a recommendation engine that feeds AI models to dynamically adjust recommendations with minimal latency, allowing retailers to personalize product suggestions based on user activity.

Real-time supply chain optimization

Kafka streams real-time data from suppliers, production lines, and inventory systems, enabling manufacturers to quickly detect disruptions, adjust production schedules, and minimize downtime, all while optimizing costs.

IoT device management

Kafka enables real-time ingestion, processing, and analysis of IoT sensor data, allowing predictive maintenance, anomaly detection, and real-time alerts needed for industrial IoT and smart city use cases.

Telecom network performance and incident management

Kafka aggregates and processes millions of network events per second to detect outages, automate incident responses, and ensure seamless connectivity. Your customers could use a network monitoring and orchestration system that helps optimize network performance and detect service outages.

Healthcare data exchange solutions

Kafka provides scalable data pipelines that allow for seamless data sharing, real-time patient monitoring, and analytics. This means your customers could securely exchange patient and clinical data with hospitals, insurers, and pharmaceutical companies in real time.


[SKIP AHEAD TO LEARN MORE >](#)

The resulting systems are highly scalable and fault tolerant, exactly what you need to build a successful product today.

Your customers are looking to you for trustworthy products that help them move faster, make smarter decisions, and operate more efficiently—all so they can outpace the competitors who are racing to the same.

As [the category leader](#), Confluent has seen the transformative potential of data streaming across industries, geographies, and business models. The use cases it unlocks are pushing companies to innovate and operate faster, raising the stakes for competition in every sector.

It can be a complex undertaking. In this ebook, **we'll give you the roadmap to building and delivering the enterprise-ready, real-time products your customers need**, whether you decide to do it yourself or accelerate your route as a Confluent partner.



The Data Streaming Report

- **86%** cite data streaming as an important or top strategic priority of IT investment in 2025
- **68%** expect the use of this technology to continue to grow over the next two years.

[READ IT HERE >](#)

1 Understanding the True Cost of Apache Kafka®

SINCE 2011, data streaming has become the default for mission-critical use cases that demand real-time data and insights. Today, Kafka is used by more than 80% of the Fortune 100.

Businesses love Kafka because it enables real-time data movement across their entire organization. With Kafka, companies can build event-driven architectures that seamlessly connect applications and data systems, ensuring they can react instantly to customer interactions, operational changes, and market shifts. Its high throughput, scalability, and fault tolerance make it ideal for handling massive data volumes, while its publish-subscribe model and stream processing capabilities allow organizations to unlock real-time insights and automation. Kafka has become the backbone of modern data infrastructures, empowering businesses to innovate faster and deliver richer customer experiences.

However, even among early adopters—some of the most innovative engineering teams in technology, telco, and financial services—**the road to operationalizing Kafka for enterprise use has never been an easy one.**

Generally, the challenges most organizations encounter along the way fall into one of two categories:

1. **Successfully operating Kafka at scale**, while meeting business requirements for reliability, security, and performance
2. **Generating value with data streaming** by powering real-time solutions for customer 360, cybersecurity, predictive analytics, IoT integration, and more

Many organizations using Kafka expend so much effort operating Kafka that they aren't able to fully dedicate time and resources to realize its full value.

DID YOU KNOW

In 2024, the global average cost of a data breach surged to \$4.88 million¹—a 10% increase year over year, according to IBM. A single security flaw or critical bug in a self-managed Kafka deployment can trigger hours—or even days—of downtime, data loss, customer fallout, reputation impact, and revenue loss. Can your business afford that risk? Or the time it would take to fix things on your own?

¹ "Cost of a Data Breach Report 2024." 2024, IBM.

Even in the early, experimentation-only phase, operating Kafka is difficult. And it only gets more complex, expensive, and time-consuming as your Kafka footprint expands. More engineers leveraging Kafka to unlock new, value-generating features is exactly what you want—but the more it happens, the more work you have to take on, including:

- **Cluster management:** Ensuring high availability and performance across distributed Kafka clusters becomes increasingly complex as the number of topics, partitions, and consumer groups grows.
- **Capacity scaling:** Proactively scaling Kafka infrastructure to meet growing data volumes and throughput demands without sacrificing performance requires substantial engineering resources and planning.
- **Data replication and consistency:** Maintaining strong data consistency across distributed systems and handling replication failures can lead to significant operational overhead.
- **Monitoring and troubleshooting:** Real-time monitoring, logging, and troubleshooting of Kafka performance, consumer lag, and broker health becomes more difficult with a growing Kafka deployment.
- **Patching and upgrading:** Regular patching and version upgrades to Kafka and its ecosystem require careful planning to avoid downtime and ensure seamless operations.
- **Security and compliance:** Implementing robust security controls such as data encryption, access control, and auditing across distributed systems becomes more resource-intensive as your Kafka ecosystem expands.

Kafka is an inherently complex technology. Managing hundreds of distributed components—like brokers, controllers, and connectors—requires precise coordination to scale with growing data volumes and throughput demands, all without sacrificing performance. Not to mention optimizing resource utilization and controlling infrastructure costs.

But Kafka's architectural complexities are also exactly what make it the ideal choice for mission-critical use cases that involve processing large volumes of data with low latency and high fault-tolerance. The faster you can get over its operational challenges, the sooner your customers can reap the benefits. That's why so many industry leaders have invested significant resources—often over multiple years—to master Kafka operations at scale and tailoring their deployments to their specific requirements.

Being able to automate Kafka deployments, operations, and management is key to minimizing downtime and optimizing operational costs. Few technical teams have the time and resources to tackle strategic, high-value work while also rebuilding Kafka's architecture for mission-critical scalability, performance, and reliability.

But remember, it's not just about mastering Kafka operations. You'll also need complementary tools to truly unlock value-generating use cases, and **capitalizing on your data streaming investment will require more than just Kafka.**

2 Kafka vs Confluent: Open Source vs Enterprise-Grade Data Streaming

SUCCESSFULLY BRINGING DATA STREAMING TO YOUR PRODUCT will require far more than just highly reliable and scalable Kafka. You need to deliver a comprehensive set of features spanning data integration, stream processing, advanced security and governance controls, monitoring, automation, support for multiple programming languages, and more.

When it comes time to scale your Kafka deployment or pursue more advanced features dependent upon the technology, you want to make sure you're equipped to succeed. Early Kafka adopters like Uber and Netflix have had to invest significant resources into infrastructure management, hiring and retaining Kafka developers and operators, and building platform capabilities around Kafka.

More Than Just Kafka



Real-time connectivity & processing

Leverage valuable legacy data to power modern, cloud-based apps & experience



Security & governance

Ensure streaming data is as safe & secure as data-at-rest as Kafka usage scales



Global availability

Maintain high availability across environments with minimal downtime

To go from running Kafka experiments to powering real-world solutions, your customers will need a complete data streaming platform—one that's ready to scale up or down on demand and has all the capabilities you need to power real-time fraud detection, IoT device management at scale, hyper-personalized digital experiences, and more.

Building these essential capabilities in-house significantly expands the scope of data streaming projects—driving up costs, increasing complexity, and delaying your time to market.

For instance, many use cases require Kafka connectors for no-code integration to common data sources and destinations, like PostgreSQL, OracleDB, Amazon S3, Google BigQuery, Snowflake, and Databricks. **But developing a single system connector can take 3-6 months of an engineering team's effort**, followed by ongoing maintenance and support throughout its lifecycle.

[SKIP AHEAD TO LEARN HOW WITH CONFLUENT >](#)

KAFKA VS CONFLUENT:

DID YOU KNOW

The Impact of a Complete Data Streaming Platform

Total Economic Impact of Confluent Platform – A 201% Return on Investment in <6 Months²



2.4 million saved

in developer and management costs



\$3.8 million gained

via accelerated business enablement

Today, you have two paths to get your data streaming service to customers:

- **Build a custom, homegrown solution:**
This requires significant investment in infrastructure, compute, and networking, along with hiring and retaining Kafka developers and operators. You'll also need to manage integrations, governance, data processing, and the ongoing complexity of scaling your Kafka deployment as your customers' needs grow.
- **Leverage a commercial distribution of Apache Kafka:**
With licensing for a complete, enterprise-ready Kafka solution—including pre-built connectors, cluster linking, and other enterprise components—you can accelerate your path to delivering a robust data streaming service without the overhead of building from scratch.

Choosing a solution with enterprise features for Kafka management and monitoring will significantly accelerate your product roadmap. But if that solution isn't a complete data streaming platform—one that has built-in features to connect, process, and govern

data streams—you'll still need to run a gap assessment and then build and maintain any missing features.

Many of the highest impact use cases we've seen customers implement—[real-time claims processing](#), [real-time payments](#), [IoT networking](#)—rely on Kafka connectors, advanced governance, and stream processing with Apache Flink®.

When combined, these tools amplify the value of data streaming, enabling your customers to quickly realize the benefits of the service you offer. Consider a real-time claims processing use case. Pre-built Kafka connectors streamline data flow between disparate healthcare systems while Flink merges, transforms, and enriches those data streams in real time to create valuable data products for downstream systems to consume and act upon. And schemas ensure consistent data standards throughout, making it easier to scale and maintain compatibility at every stage.

²"The Total Economic Impact™ of Confluent Platform," 2018. Forrester.

7 Features of a Complete Data Streaming Platform

When evaluating data streaming solutions, you need a platform that goes beyond core streaming features of Kafka.

Look for These 7 Features & Capabilities

- ✓ Multi-language development tooling
- ✓ DevOps tools for management and monitoring
- ✓ Autobalancing, data replication, and disaster recovery features
- ✓ Low-cost and low-latency storage configurations
- ✓ Pre-built, expert-verified connectors to popular data sources and sinks
- ✓ Enterprise security and data quality controls
- ✓ Advanced stream processing

With these capabilities at your disposal, your organization can power unlimited use cases and better business outcomes for your customers—and new revenue streams. So are you ready to dedicate time and resources to building advanced platform capabilities around self-managed Kafka?

Confluent’s mission is to unlock everything data can do for the modern organization, and it all starts with data streaming.

Our founders, the original co-creators of Apache Kafka®, established a team of the leading experts in data streaming and stream processing. In the decade since, we’ve built a complete platform that’s been adopted by industry leaders and validated by technology analysts. Confluent provides the best data streaming experience, in the cloud, on-premises, and at the edge.

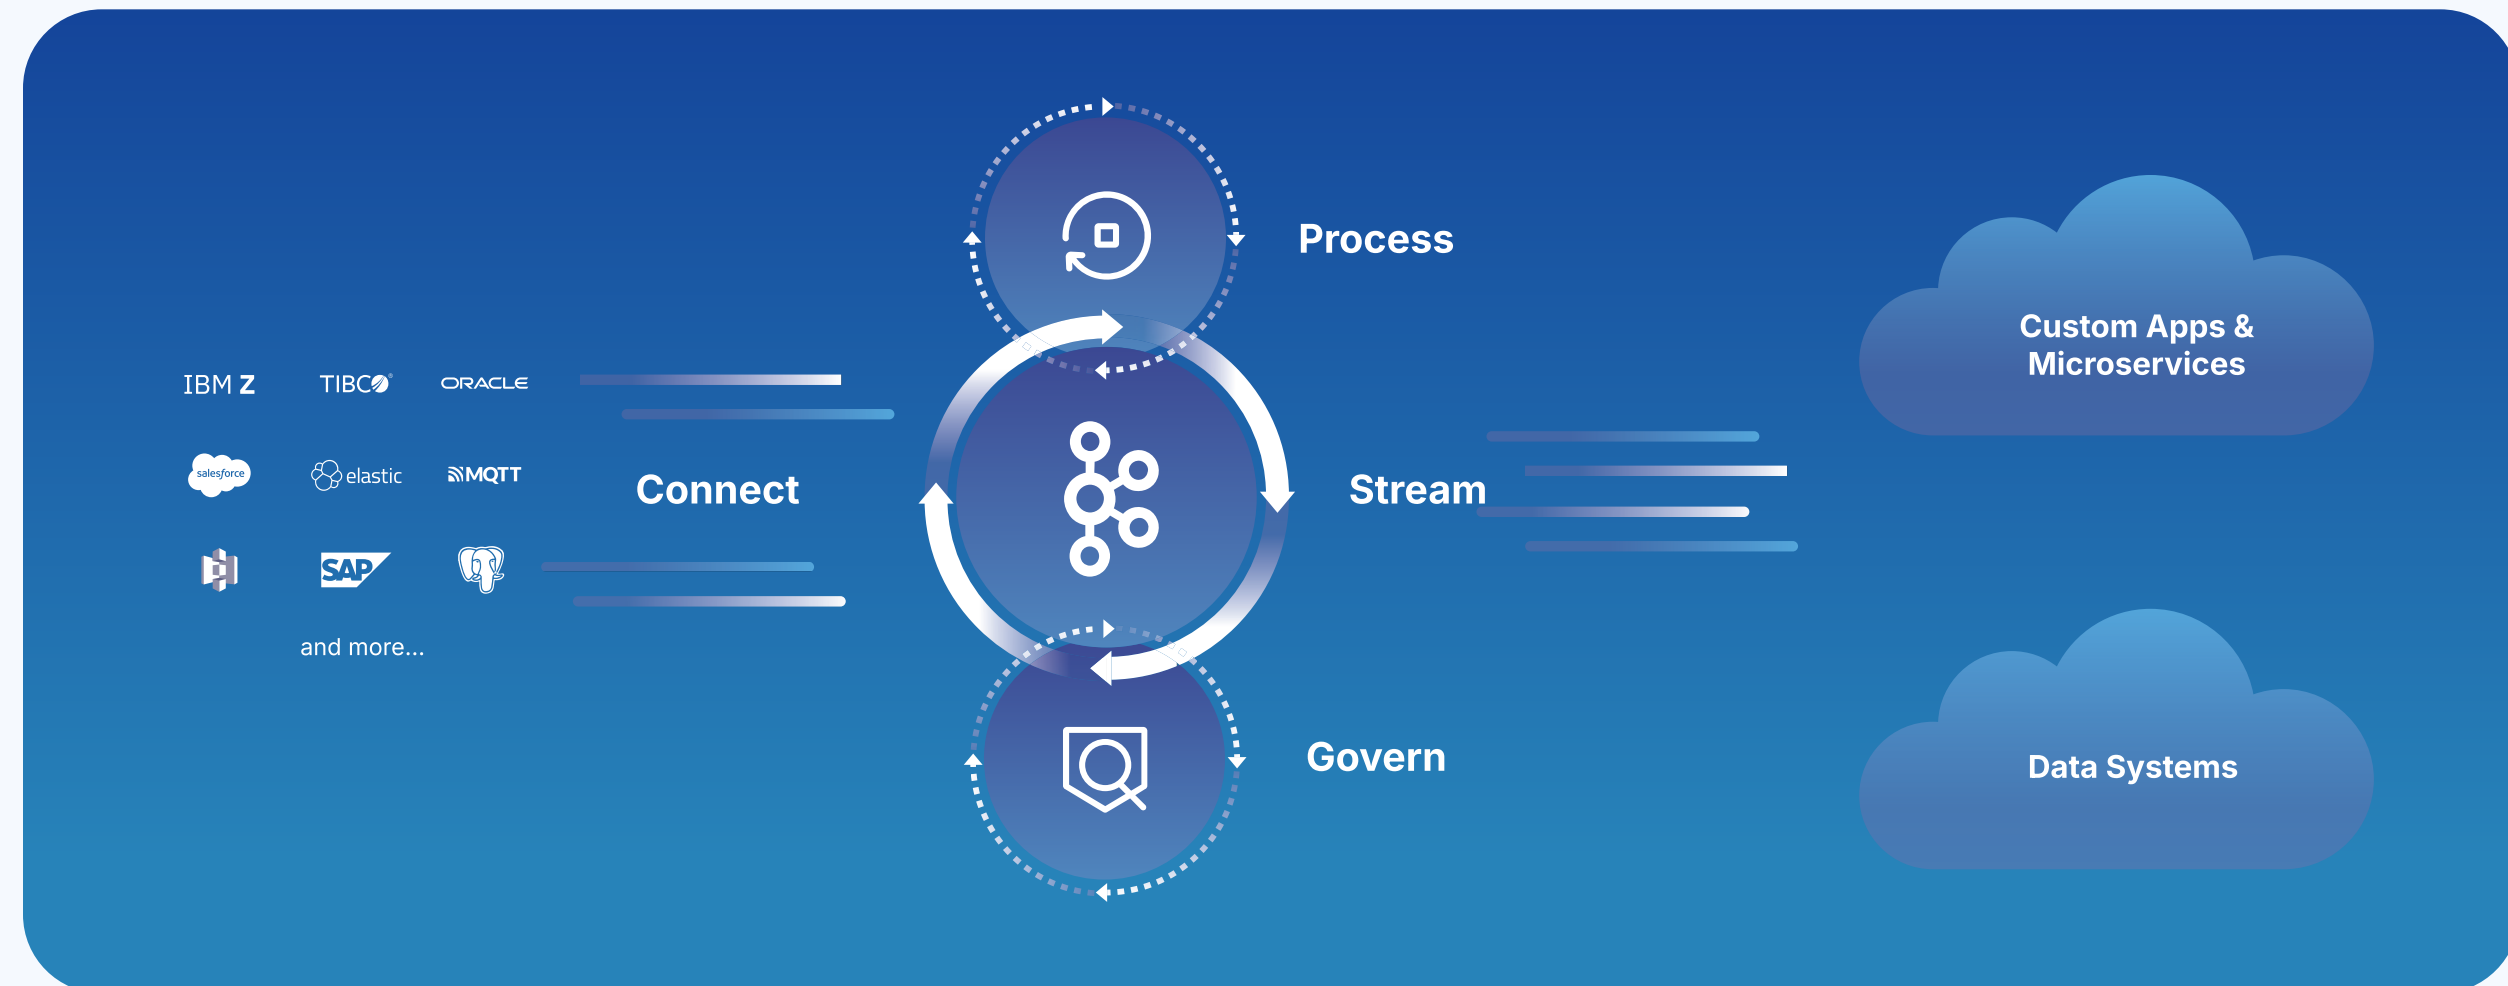
Keep reading for a deep dive into our self-managed and BYOC deployments, Confluent Platform and WarpStream.

A Complete Data Streaming Platforms Unlocks Unlimited Use Cases in...

RETAIL	Inventory Management	Personalized Promos	Product Development	Sentiment Analysis	Streaming Enterprise Analysis	Systems of Scale for High Traffic Periods
FINANCE	Early-On Fraud Detection	Capital Management	Market Risk Recognition & Investigation	Preventive Regulatory	Real-Time What-If Analysis	Trade Flow Monitoring
TRANSPORT	Advance Navigation	Environmental Factor Processing	Fleet Management	Predictive Maintenance	Threat Detection & Real-Time Response	Traffic Distribution Optimization
TELECOM	5G Networks	Billing Reconciliation	Product Development & Introduction	Sentiment Analysis	IOT Integration	Systems of Scale for High Traffic Periods
ALL INDUSTRIES	Data Pipelines	Hybrid Cloud Integration	Microservices	Security & Fraud	Customer 360	Streaming ETL

Meet the Confluent Data Streaming Platform

Stream, connect, process, and govern your data with an all-in-one, real-time platform from the pioneer in data streaming.



3 A Deep Dive

Confluent Platform and WarpStream

A DEEP DIVE

YOU CAN BUILD, LAUNCH, AND GROW FASTER wby embedding Confluent’s data streaming platform into your business. No matter where your business is run, we provide flexible deployment options, distribution licensing, and competitive commercial terms through our OEM program.

[SKIP AHEAD TO LEARN MORE >](#)

Before we go over the program, let’s dive into the details of our software and bring-your-own-cloud (BYOC) deployment options: [Confluent Platform](#) and [WarpStream](#).

Confluent Platform

At Confluent, running Kafka at scale and in production is what we do better than anyone else. Today, we successfully manage more than 30,000 Kafka clusters for our customers, with 3 trillion messages written per day and more than 1 exabyte of data processed per year.

We’ve applied that technical expertise to completely rearchitect the open source engine and deliver an enterprise-grade software offering, Confluent Platform.

On its own, the open source distribution of Kafka is an incredibly powerful technology, powering some of the highest-throughput platforms in the world. But if data streaming and processing are critical capabilities for you or your customers’ needs, you need to hit the ground running to keep up with the competition.

Confluent Platform includes all the tools and capabilities to deliver the comprehensive platform today’s enterprise organizations need for their most mission-critical data streaming projects.

GOOD TO KNOW

Altogether, Confluent reduces the total cost of ownership (TCO) for Kafka deployments by up to 40%³ so both you and your customers can reinvest valuable resources back into innovation and true differentiation.

³“Measuring the Cost-Effectiveness of Confluent Platform,” 2021. Confluent.

Stream

Confluent brings a highly reliable, enterprise-grade Kafka experience to any environment where data streaming is needed—at any scale.

- Fluctuating, unforeseen demand is easily managed [with elastically scaling](#) clusters that automate partition rebalances.
- With [tiered storage](#), infinite amounts of data can be stored right within Kafka while cost-effectively separating storage from compute.
- Downtime costs and business disruption are minimized with clusters deployed across [multiple regions](#).
- And finally, data streams can be automatically synced wherever they are needed—in the cloud, across clouds, etc.—with [Cluster Linking](#).

Connect

With [pre-built connectors](#), Confluent Platform ensures organizations can leverage data streaming across the entire business. Avoid 3-6 engineering months of design, build, and test time required for each connector while ensuring your customers can:

- Quickly connect data systems and apps leveraging a rich ecosystem of 120+ connectors built by Kafka experts.
- Boost developer productivity with built-in connector usability and data transformation features.
- Easily build streaming data pipelines using connectors to bridge legacy systems to modern, cloud technologies.

Process

Confidently unify data streaming, governance, and on-the-fly processing with [Confluent Platform for Apache Flink®](#):

- Enables high-performance and efficient stream processing at any scale
- Filter, analyze, and enrich data streams
- Power low-latency applications and pipelines
- Unlock advanced real-time and event-driven use cases

Govern

With [Schema Registry](#), your customers will be able to reduce operational complexity:

- Standardizing on schemas with stored version history
- Validating data compatibility at the client level
- Removing time-consuming coordination among developers
- Syncing schemas across Kafka clusters in real time
- Maintaining trustworthy, high-quality data streaming across environments

Comparing DIY Kafka and Confluent Platform

	Open Source Kafka	Confluent	
Distributed Event Streaming	✓	✓	Core Data Streaming & Processing
Kafka Connect Integration Framework	✓	✓	
Kafka Streams Client Library	✓	✓	
Flexible Development	-	Non-Java Clients, Admin REST APIs	Enterprise Scalability
Elasticity	-	Self-balancing Clusters	
Cost-effective Storage	-	Tiered Storage	
High Availability	-	Multi-region Clusters	
Pre-built Connectors	-	120+ Connectors	
Advanced Stream Processing	-	Apache Flink®	
Data Quality Controls	-	Schema Registry, Schema Validation	
Security Controls	-	RBAC, Audit Logs, Secret Protection	
Global Resilience	-	Cluster Linking, Replicator	Enterprise Management
Automation	-	Confluent for Kubernetes, Ansible Playbooks	
Monitoring	-	Control Center, Health+	



“Businesses undergoing digital transformations rely on Infosys for technologies that meet rigorous enterprise requirements and ensure long-term success. By leveraging Confluent’s advanced data streaming solutions, we can complete customer projects more efficiently, with minimized risk and reduced costs.”

Dinesh Rao, EVP and Co-Head of Delivery, Infosys

Secure

Ensuring [confidentiality, compliance and privacy](#) always remains a top priority. Our customers and partners depend on our enterprise-grade security features including:

- Secret Protection
- Structured Audit Logs
- Role-based access controls (RBAC)
- OAuth / OIDC support

Enterprise-Grade Security



Secure Protection

Secret Protection safeguards all critically sensitive information within Kafka with at-rest encryption of configuration files.



Structured Audit Logs

Structured Audit Logs captures authorization logs in a set of dedicated Kafka topics, on a local or a remote cluster.



Role-Based Access Controls (RBAC)

RBAC is a centralized implementation for secure access to Kafka resources with fine-tuned granularity and platform-wide standardization.



OAuth / OIDC Support

Streamline authentication by managing application identities and credentials through your own OIDC identity provider with OAuth, an industry standard for authentication.

Monitor and Automate

We provide advanced monitoring and automation capabilities to ensure seamless data streaming operations within both production and development environments.

- With Health+ Monitoring, businesses can identify and prevent cluster outages through intelligent alerting made available through cloud-based monitoring tools.
- For efficient management, Confluent for Kubernetes provides a declarative API to automate and streamline operations for deployment to any standard or managed Kubernetes environment.
- Ansible playbooks enable automated deployments in non-containerized, virtual, or bare metal environments, ensuring smooth and reliable performance wherever data streaming is needed.

With Confluent Platform, our goal is to make it as easy as possible to manage Kafka workloads. This solution is one that many of our partners use to deliver seamless, enterprise-grade data streaming to their own customers, without having to shift their attention away from their core competencies and engineering priorities.

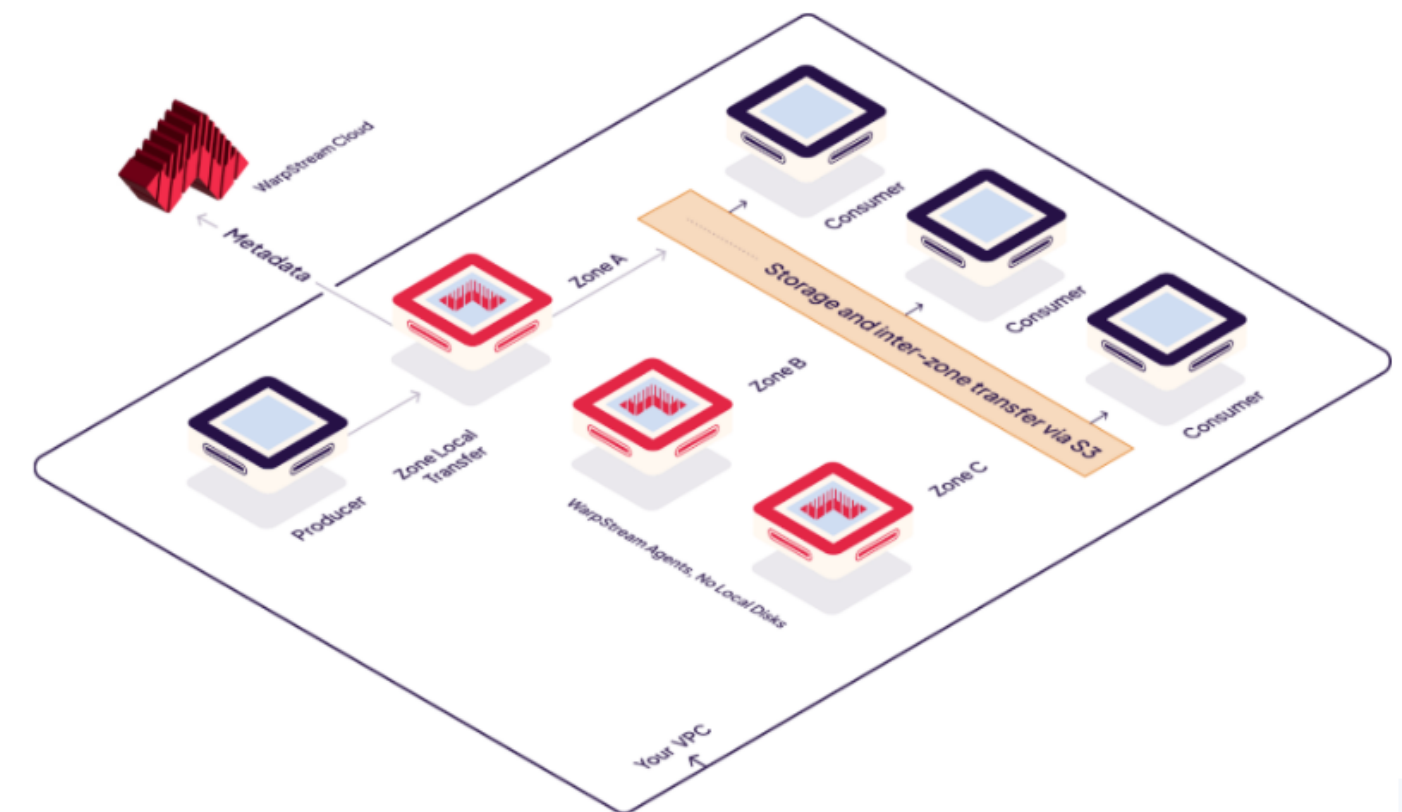
WarpStream

WarpStream is Confluent's bring-your-own-cloud (BYOC) deployment option, ideal for organizations with stringent security and data sovereignty requirements, particularly in finance, healthcare and the public sector.

WarpStream delivers:

- **Zero-access security by default.** WarpStream operates on a BYOC-native deployment model, so your data never leaves your environment.
- **Low cost networking for relaxed latency scenarios.** Perfect for “firehouse workloads” like logging, monitoring, telemetry, and feeding data lakes.
- **Kafka-compatible data streaming with any Kafka client,** and identical Kafka nomenclature such as topics, partitions, and groups.

With WarpStream, there are no inter-AZ networking costs and no disks to manage—just infinite scalability, all within your virtual private cloud.



LEARN MORE ABOUT WARPSTREAM >

Have questions about whether Confluent Platform or Warpstream is right for your needs?

[Get in touch with our team](#)

4 Confluent's OEM Program

Grow Faster With Enterprise-Grade Data Streaming

BRING DATA STREAMING TO YOUR BUSINESS quickly and confidently with unified Apache Kafka® and Apache Flink®—built and supported by the original co-creators of open source Kafka.

OEM partners receive licensing to embed [Confluent Platform](#) and [WarpStream](#) within their service offerings.



Monetize data streaming together with the industry leader and unlock more real-time customer use cases on premises, at the edge, and in the cloud



Maintain engineering focus on your core roadmap and accelerate launch plans with the industry's most complete, ready-to use data streaming platform



Reduce risks and ensure customer success with implementation guidance, certification, flexible commercial terms, and technical support from the Kafka experts

The OEM Program gets you to market faster—allowing you to build, launch, and operate confidently, backed by the support and guidance of the world's foremost experts in data streaming technology. The program features:

- **Design Reviews and Development Support:** Build with architectural guidance and training from the team with 5 million hours of Kafka development logged.
- **Technology Certifications:** Launch confidently with publicized backing from the data streaming leader.
- **Flexible Commercial Terms:** Package your offering easily with commercial terms that match the way you sell.
- **Enterprise Support:** Bring expert support to your business and easily handle any customer question or issue.

With the OEM program, you'll have peace of mind that you're equipped to meet your customers' needs. **Upon joining, you gain access to an exclusive partner portal with on-demand training and other go-to-market enablement resources that will allow your team to build and launch faster, with confidence.**

The program allows you to bypass the complexities, expenses, and risks of open source projects and operate an enterprise-grade service with step-by-step guidance from the data streaming experts.

From the conceptualization of your new offering to service launch, we will help you to decide the best way to build and deploy your data streaming platform. Our team has millions of hours of Kafka development and operations logged—all of that expertise will be available for you as a Confluent partner. We will help you evaluate and minimize technical tradeoffs based on your business goals and ensure you fully understand how to leverage the data streaming platform for your customers.

ARE YOU READY?

Ready to Jump Start Your Data Streaming Journey?

With Confluent, you don't just get immediate access to a complete data streaming platform. You also gain expert guidance across the entire data streaming lifecycle.

[Contact us to learn more](#)

Next Steps

Start Building Real-Time Products With Enterprise-Grade Data Streaming

START BUILDING REAL-TIME PRODUCTS WITH ENTERPRISE-GRADE DATA STREAMING—one that streams, connects, processes, and governs data wherever it resides—you can build the products and data services your customer need.

Focus your resources and in-house expertise on what matters most: your core business. When it comes to data streaming, we'll handle the rest.

Get in Touch

Get in touch with our team of Kafka experts today and learn how Confluent can help you quickly advance your engineering roadmap.

[GET IN TOUCH WITH OUR TEAM >](#)



CONFLUENT

Want to see Confluent in action for yourself first? You can try Confluent Platform for free—30 days, unlimited brokers.

[Download for Free >](#)

[Explore the Quick Start >](#)

[Follow a Demo and Guided Tutorial >](#)



WarpStream

Start your free trial of WarpStream. Get \$400 in free credits, no credit card required.

[Try WarpStream for Free >](#)



Additional Resources

[Measuring the Cost Effectiveness of Confluent Platform](#)

[Introducing Confluent OEM Partner Program](#)

[Apache Kafka vs Confluent](#)

[2025 Data Streaming Report](#)

