CONFLUENT

EBOOK

Delivering Data Streaming-as-a-Service:

THE ENTERPRISE GUIDE

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

Contents

Confluent's OEM Program

Grow Faster With Enterprise-Grade Data Streaming

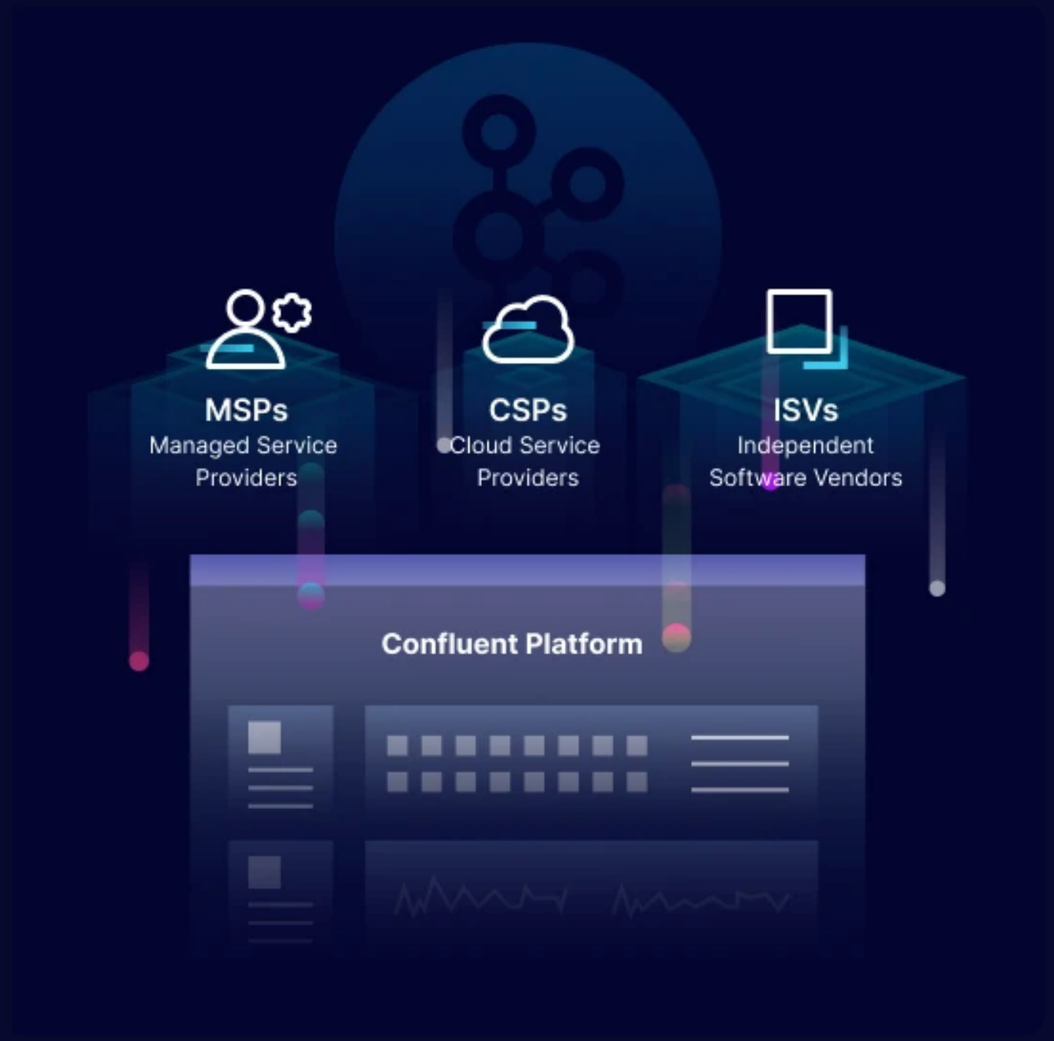

Executive Summary

DATA STREAMING IS THE KEY TO UNLOCKING REAL-TIME ANALYTICS, modern application development, and data-driven operations—not to mention operationalizing artificial intelligence (AI) and so much more. And with more than 150,000 companies having adopted the technology, Apache Kafka® has become the de facto standard for making "real time" a reality.

Your customers need data streaming to remain relevant and stay ahead of their competitors—but can you afford to build and manage the Kafka service they expect all on your own? With all the requests hitting your business and the rapid pace at which AI is transforming markets, the development time involved poses a significant opportunity cost.

While immensely impactful, Kafka's distributed architecture makes it inherently complex and difficult to operate. However, leveraged correctly, it is the ideal choice for processing large volumes of data with low latency and high fault-tolerance.

When evaluating data streaming solutions, your customers will need a platform that goes beyond the core features of open source Kafka. One that allows them to get data streaming projects and use cases out the door faster.

Built by the original co-creators of Kafka, Confluent is the indisputable leader in data streaming, so we know exactly what it will take to get a production-ready, enterprise-grade data streaming service to market that your customers will love. Our data streaming platform operates more than 30,000 Kafka clusters every single day on behalf of the largest, most innovative companies throughout the world.

We'll show you the fast-track to delivering a world-class data streaming service: Confluent's OEM Program. As an OEM partner, you'll leverage Confluent Platform and WarpStream to bring a complete data streaming solution to market faster with minimal investment, risk, or distraction from your core business. Your customers can stream, connect, process, and govern all their data in real time—backed by world-class training, support, and guidance from the leading experts in data streaming.

Ready to launch your data streaming service?

LET'S GET STARTED.

SKIP AHEAD TO LEARN ABOUT CONFLUENT'S OEM PROGRAM >

Confluent Powers Data Streaming Across Industries

Alibaba Cloud

"Customer demand for data streaming has skyrocketed as businesses strive for a competitive edge through application modernization, real-time analytics, and AI. With Confluent, we were able to deliver a complete, enterprise-grade Apache Kafka managed service on Alibaba Cloud with minimal time and engineering effort."

Dongliang Guo, Vice President of International Business, Head of International Products and Solutions, Alibaba Cloud Intelligence

Infosys®

"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

Introduction

Meeting the Demand for Data Streaming

MODERN BUSINESSES RISE AND FALL BASED ON THEIR DIGITAL CAPABILITIES. Software has transformed everything, and organizations no longer try to keep every technical skill set in-house. Instead, they focus on their core business and turn to experts for specialized technologies.

That's why enterprises routinely rely on managed and cloud service providers (MSPs, CSPs) for critical IT functions—from security and infrastructure management to IT support and disaster recovery. These services help them control costs, ensure seamless interoperability across their tech stack, and reduce risk.

The more challenges you solve for your customers, the more they'll trust you to deliver the cutting-edge solutions they need to succeed. **And now, there's a new challenge at hand: helping customers harness data streaming to modernize application development, operate with real-time business insights, and unleash the full power of AI.**

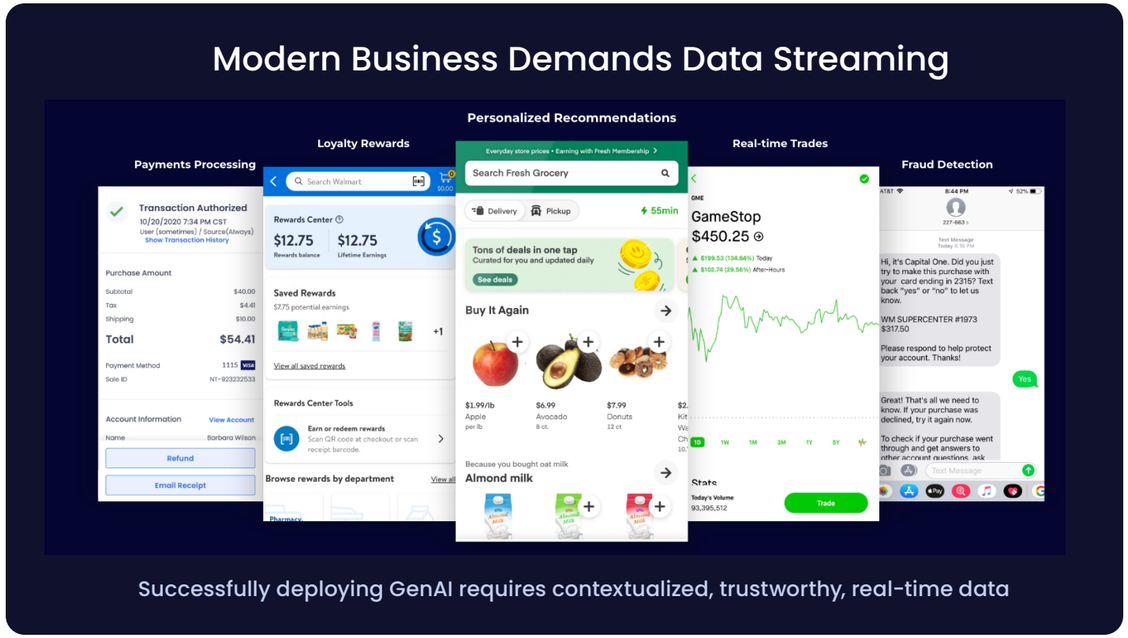

DID YOU KNOW

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.

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.

When implemented correctly, Kafka empowers organizations to build microservices architectures, enhance operational efficiency, and support agile, data-driven decision-making—all key drivers of digital transformation. 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.

Consider how your customers can leverage Kafka for:

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 data streaming service to build 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 develop solutions that enable them to securely exchange patient and clinical data with hospitals, insurers, and pharmaceutical companies in real time.

SKIP AHEAD TO LEARN MORE >

Building and managing data streaming capabilities will help your customers unlock these highly valuable use cases and find more overall success with your business, but it's a complex and expensive undertaking.

In this ebook, we'll give you the roadmap to delivering an enterprise-ready data streaming service, whether you decide to do it yourself or accelerate your path as a Confluent partner.

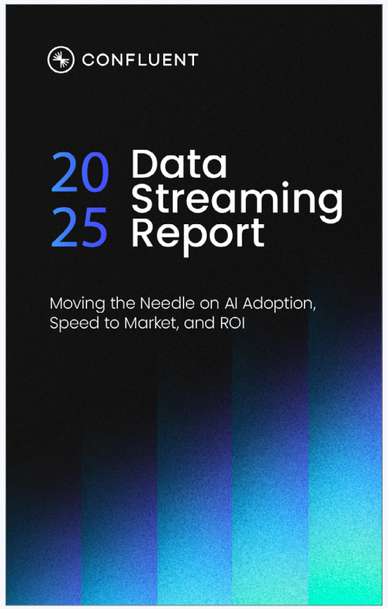

The Data Streaming Report

- **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.**

Even in the early days, operating Kafka can be difficult. And it will only get more complex, expensive, and time-consuming as you onboard more customers and your Kafka footprint expands. More customers leveraging Kafka to unlock new, value-generating use cases is exactly what you want—but the more it happens, the more work you have to take on, including:

- **Cluster management and capacity scaling:** Ensuring high availability and performance across distributed Kafka clusters becomes increasingly complex as the number of topics, partitions, and consumer groups grows. This complexity becomes even worse without the right tools and expertise in place.
- **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.

DIY Kafka in Action

To unlock the value of data streaming for your customers, you'll need to operationalize Kafka for massive scale and build a complete set of platform capabilities around it. Consider the system architecture of this early stage data streaming platform.

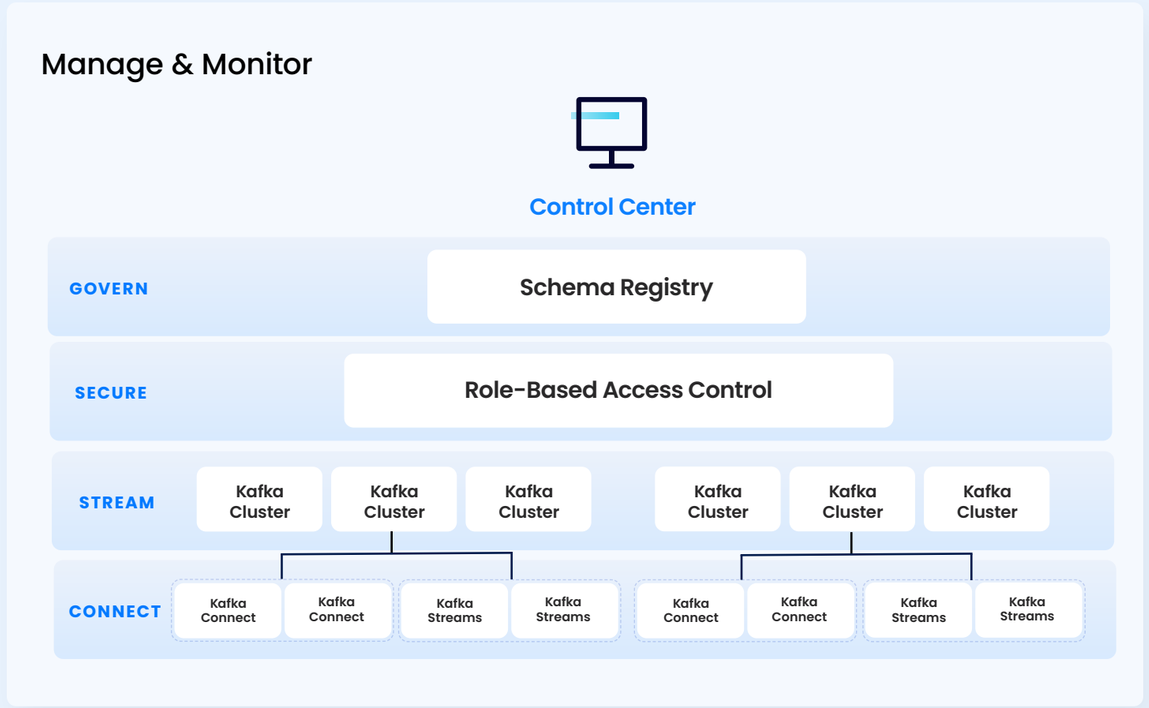

Our on-premises and private cloud distributions, Confluent Platform and WarpStream, let you skip the DIY and get to market faster.

DID YOU KNOW

In 2024, the global average cost of a data breach surged to $4.88 million[1]—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?

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 help customers get over its operational challenges, the sooner they can reap all the benefits of having real-time data streaming and processing leveraged throughout your business and platform.

Being able to automate Kafka deployments, operations, and management is key to minimizing downtime and maximizing value for your customers. This allows them to get new projects and use cases out the door faster, while optimizing operational costs for reinvestment into net-new projects. Few technical teams will have the time and resources available to tackle strategic, high-value work while simultaneously managing Kafka for mission-critical scalability, performance, and reliability.

But remember, it's not just about mastering Kafka operations. There's a wide set of tools your customers will need to unlock value-generating use cases, and **capitalizing on the demand for data streaming requires more than just Kafka.**

Solve These 6 Challenges to Unlock Kafka's Full Potential

1. Lack of expertise and resources
2. Difficulty moving from development to production
3. Unpredictable outages and downtime
4. Difficulty securing streaming data
5. Lack of governance
6. Difficulty scaling

READ THE EBOOK: THE TOP 6 REASONS KAFKA PROJECTS FAIL >

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

2 Spot the Difference

Open Source vs Enterprise-Grade Data Streaming

SUCCESSFULLY DELIVERING A DATA STREAMING SERVICE requires 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 for customers to scale their Kafka deployment or pursue more advanced use cases, you want to make sure they're equipped with everything necessary to succeed.

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 needed to power real-time analytics, fraud detection, AI/ML, hyper-personalized digital experiences, and much 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, most use cases will require Kafka to be connected to common data sources and destinations, like PostgresSQL, 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 >

DID YOU KNOW

The Impact of a Complete Data Streaming Platform

Total Economic Impact of Confluent Platform - A 201% Return on Investment in <6 Months[2]

2.4 million saved
in developer and management costs

$3.8 million gained
via accelerated business enablement

aws Financial Services Competency 2023 · Google Cloud Partner of the Year 2023 · Microsoft US Partner of the Year 2023 · mongoDB. Partner of the Year 2022 · FORRESTER® WAVE LEADER 2023 Streaming Data Platforms · JPMorgan Chase & Co. Strategic Partner of the Year 2022 · Morgan Stanley CTO Innovation 2019 · Bank of America Enterprise Technology Innovation 2019

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.

Licensing 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.

[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

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 deployment options, Confluent Platform and WarpStream.

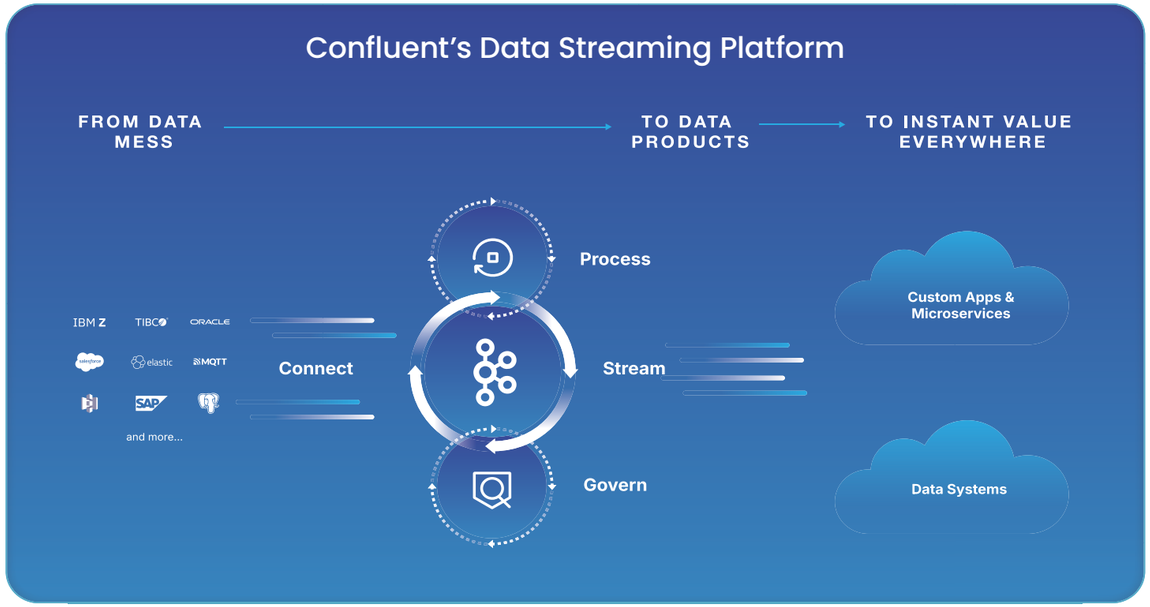

Confluent's Complete Data Streaming Platform Unlocks Unlimited Use Cases in...

RETAIL	Inventory Management	Personalized Promos	Product Develoment	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

3 A Deep Dive

Confluent Platform and WarpStream

YOU CAN EMBED CONFLUENT'S DATA STREAMING PLATFORM wherever your business needs it—in the cloud, on-premises, or at the edge—with distribution licensing available 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 Stream Governance, 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	
Elasticity	-	Self-balancing Clusters	
Cost-effective Storage	-	Tiered Storage	
High Availability	-	Multi-region Clusters	**Enterprise Scalability**
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	
Automation	-	Confluent for Kubernetes, Ansible Playbooks	**Enterprise Management**
Monitoring	-	Control Center, Health+	

MINDGATE™
DELIVERING INNOVATION

"Processing over 7 billion real-time payment transactions per month for leading financial institutions worldwide comes with exceptional performance expectations. **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.

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.

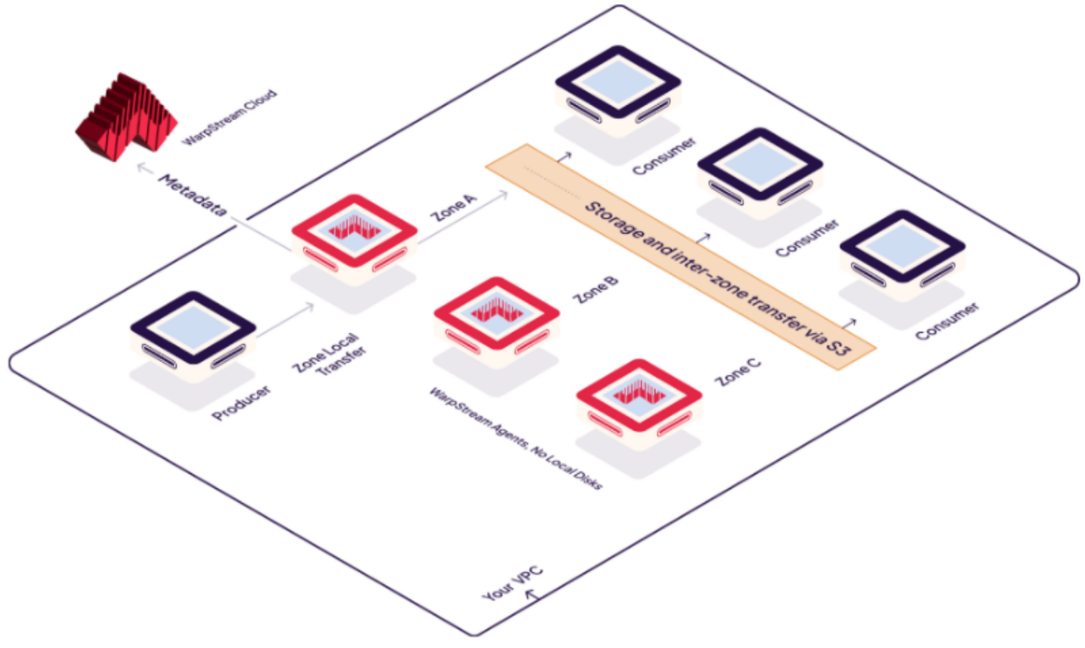

LEARN MORE ABOUT WARPSTREAM >

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 CUSTOMERS 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

Deliver Enterprise-Ready Data Streaming-as-a-Service

DATA STREAMING IS IN HIGHER DEMAND THAN EVER but can come with a complex, high price tag that introduces unnecessary risk to a business. With easy access to a complete data streaming platform, you can deliver the Kafka service your customers need—one that allows them to stream, connect, process, and govern all of their data at scale, wherever it resides.

Becoming an OEM partner and building a business with Confluent accelerates your time to market. And it unlocks enormous benefits—limitless use cases for your customers, new revenues for you, and the opportunity to deepen those relationships as you solve new challenges together.

Don't wait years to build, optimize, and deliver the service your customers need now.

Get in Touch

Get in touch with our team of Kafka experts today and learn how Confluent can help you quickly launch the most successful data streaming service.

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

Leveraging Confluent Platform and OEM Program >

Measuring the Cost Effectiveness of Confluent Platform >

Webinar: Introducing Confluent OEM Partner Program >

Apache Kafka vs Confluent >

2025 Data Streaming Report >

CONFLUENT