

UNPACKING THE TRUE COSTS OF OPEN SOURCE KAFKA AND MSK

Good Teams Manage Apache Kafka®. Efficient Teams Use Confluent.

Introduction	2
Confluent's Approach to Measuring Cost-Effectiveness.	2
Confluent Versus Self-Supporting Open Source Kafka.	3
The Direct and Indirect Costs of Self-Supporting Apache Kafka	3
Reduce Data Streaming TCO by Up to 60%.	3
Reduced Infrastructure Footprint	4
Efficient Operations At Scale.	5
Minimized Risk and Downtime	5
Accelerate Time-to-Value from Months to Minutes	6
Maximize Your ROI with a 3X Return on Your Kafka Spend.	7
Confluent Cloud versus MSK	8
Infrastructure	9
Operational Costs.	9
Intangible Costs	9
Conclusion.	10
About Confluent	10

Introduction

At Confluent, we believe data in motion is at the center of the next major wave of technology innovation to enable the modern business. Data in motion enables organizations to connect, process, and react to their data streams in real time to power applications and operations across their business. Rather than passively storing data at rest and periodically querying it, modern enterprises are building rich, digital customer experiences and automated business operations by continuously processing data streams in real time.

Used by over 70% of the Fortune 500, Apache Kafka® has emerged as the de facto standard infrastructure for setting data in motion and powering these modern digital initiatives. As companies accelerate their efforts to compete in today's dynamic, digital-first landscape, Kafka has become mission-critical infrastructure for the modern data stack and essential to business survival.

However, self-supporting Kafka in-house incurs significant costs in terms of infrastructure footprint, engineering resources required, and the risk of business disruption. As teams navigate through periods of economic uncertainty and the associated constraints, they are challenged with deploying and scaling infrastructure and managing day-to-day operations with reduced budgets and headcount. Some teams turn to other cloud-hosted Kafka services in an effort to reduce costs and operational burden, but simply hosting data systems on cloud infrastructure is not enough to realize the full value and benefits of the cloud.

That's why so many companies turn to Confluent to help reduce their data streaming total cost of ownership (TCO) by up to 60%, accelerate time-to-market 6+ months faster, and deliver over 3X return on investment (ROI)—all while powering a scalable and future-proof data backbone across their entire business. When teams are challenged to "do more with less," only Confluent provides a truly cloud-native and complete Kafka service that is 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.

This whitepaper outlines our approach to calculating the cost-effectiveness of our data streaming platform. We outline the business value assessment model we use, the consultative approach we take with our customers, and lessons we've learned in the process of completing real-world assessments across a full range of industries and maturities. To demonstrate the business value of using Confluent, we also outline some examples comparing:

- Confluent vs. self-managed open source Kafka
- Confluent Cloud vs. Amazon MSK

Confluent's Approach to Measuring Cost-Effectiveness

In order to illustrate the full business value of our data streaming platform, we look at three value buckets:

- **Reduce total cost of ownership**—Operate more efficiently with lower infrastructure cost, maintenance, and downtime risk
- **Increase time to value**—Deploy Kafka at scale within one week of starting with Confluent
- **Maximize return on investment**—Deliver higher returns with your project by launching faster and reducing operational burden

Reducing costs and operating efficiently is always a top priority for businesses, especially for the data architecture that underpins their digital operations. When evaluating the TCO value bucket, we look at the costs of managing a complex distributed system at scale. These costs can be organized into three broad categories:

- 1. Infrastructure costs** include the compute, storage, and networking resources allocated to provisioned clusters. The cloud has made infrastructure utilization more efficient, but the costs can still add up as new data in motion use cases are implemented, additional clusters are spun up, and more and more data is ingested. Deploying a self-managed option on premises often exacerbates these costs.
- 2. Operational costs** capture the time, engineering resources, and third-party services needed for deploying and managing a data streaming platform. In addition to the engineers required to build and maintain tooling to deploy safely and reliably to production, this category also encompasses the engineers needed to operate the platform on a day-to-day basis. This can also include third-party spend to cover professional services, training, support plans, etc.
- 3. Intangible costs** and business risk include the costs associated with unexpected cluster downtime and security incidents. While these costs can be difficult to quantify, their significance becomes abruptly apparent when an incident does occur. In addition, valuable resources have to be diverted to address unplanned downtime and breaches.

Confluent Versus Self-Supporting Open Source Kafka

The Direct and Indirect Costs of Self-Supporting Apache Kafka

Although Kafka has become the de facto standard for data streaming, managing and scaling Kafka in-house brings unnecessary burden, expense, and risk. Some of the key challenges of running Kafka on your own include:

- **Operational burden and resources**—Managing and scaling the platform to support ever-growing demand
- **Security and governance**—Ensuring streaming data is as safe and secure as data-at-rest as Kafka usage increases
- **Real-time connectivity and processing**—Leveraging valuable legacy data to power modern, cloud-based apps and experiences
- **Global availability**—Maintaining high availability across environments with minimal downtime

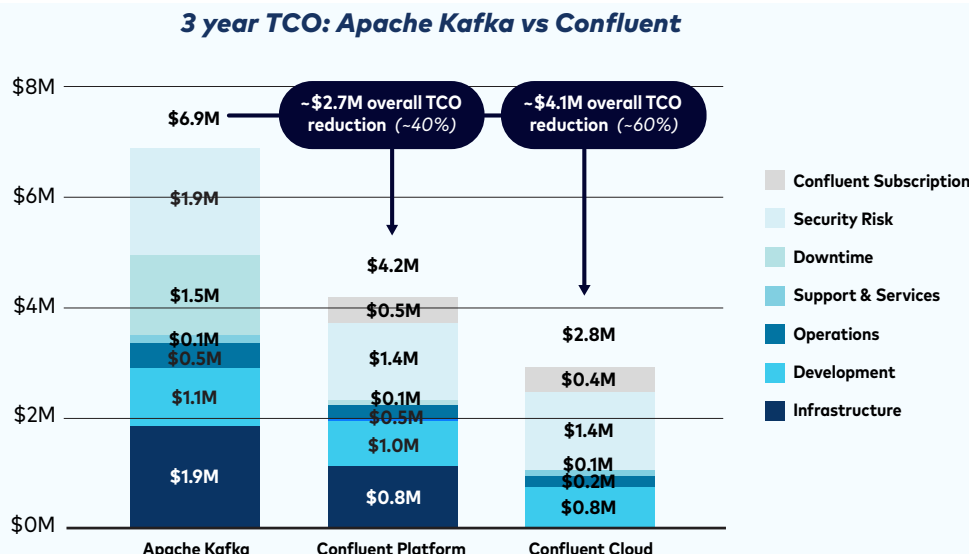
These challenges only get harder as your Kafka footprint grows across increasingly mission-critical data and use cases throughout your organization. What impact does this have on your business? Kafka may be free to download, modify, and use, but self-supporting it in-house means that you also bear the substantial direct and indirect costs in-house as well:

1. In terms of infrastructure costs, over-provisioning capacity to support fluctuating demand leads to **hundreds of thousands of dollars spent annually on underutilized and sub-optimal infrastructure**. With Apache Kafka, storage and compute are tied together. There is a practical limit to how much you can store on a single Apache Kafka broker. Once that limit is hit, you have to provision additional brokers and pay for more resources than otherwise necessary.
2. When considering operational costs, your valuable and talented engineers are spending their time building and maintaining low level infrastructure tooling instead of on more strategic projects that differentiate your business. It's also increasingly difficult to hire and retain Kafka talent because it requires a highly sought after and specialized skill set. This can result in **~3-5M+ in ongoing platform development and operations costs**.
3. Finally, in terms of business risk, it's critical to ensure your Kafka use cases are resilient and secure to avoid unwanted downtime and breaches. This risk only increases as Kafka spans more use cases, apps and data systems, teams, and environments across your organization. Downtime and security breaches have latent long-term consequences, including **reputational damage, reduced customer satisfaction, revenue and data loss, and other negative outcomes resulting from a lack of regulatory compliance**.

Let's see how Confluent measures up against self-managed Apache Kafka.

Reduce Data Streaming TCO by Up to 60%

For a customer self-supporting open source Kafka, we first try to estimate the cost breakdown of a customer's existing operating expenses and TCO cost drivers. When reviewing cost optimization opportunities in close partnership with our customers, we typically see something like this:



Confluent provides a number of essential features that make data streaming vastly more cost-effective across all the major cost categories. This results in overall TCO savings of up to 60% with Confluent Cloud and 40% with Confluent Platform.

Confluent Feature	Reduced Infrastructure Footprint	Efficient Operations at Scale	Minimized Risk and Downtime
Cloud-native Performance & Elasticity			
Serverless	●	●	●
Elastic Scaling	●	●	●
Infinite / Tiered Storage	●	●	
Self-Balancing Clusters	●	●	●
Stream Processing and Integration			
120+ Pre-built Connectors	●	●	●
Flink	●	●	
Global Resilience			
Cluster Linking	●	●	●
Multi-Region Clusters	●	●	●
Committer-Driven Expertise	●	●	●
99.99% SLA		●	●
Enterprise Security & Governance			
RBAC, Audit Logs, & Secret Protection	●	●	●
Stream Governance		●	●
Management & Monitoring			
Cloud UI / Control Center		●	●
Health+	●	●	●

Reduced Infrastructure Footprint

Self-managing Kafka on private infrastructure means considerable infrastructure costs, including compute, storage, and networking. Running hosted Kafka in the cloud still means potentially significant cloud provider bills. For example, if you are going to run Kafka on AWS, you'll need to pay for EC2 machines to run your brokers—machines that you're on the hook to spin up and down on your own. If you are using a Kubernetes service like EKS, you pay for nodes and for the service itself (e.g., Kubernetes masters). On top of this, there are network costs. Getting data into EC2 costs money, and depending on your network setup (e.g., VPC, private link, or public internet), you may need to pay both when sending and receiving data.

Confluent Cloud completely eliminates private infrastructure costs with a fully-managed Kafka service that's been completely re-architected to harness the full power of the cloud. A serverless architecture means your development teams can build and run applications without having to worry about the underlying infrastructure. You can also elastically scale up to meet any data in motion workload and scale it back down to avoid over-provisioning. Our usage-based model means you only pay for what you use and the underlying infrastructure dynamically scales as needed without interrupting running applications.

In terms of storage, Kafka platform operators have traditionally needed to trade off between offering their developers longer data retention periods and growing infrastructure costs—with open source Kafka, the longer the retention period, the more compute and storage hardware that is required. This is no longer a trade-off Kafka platform operators need to make with Infinite Storage in Confluent Cloud. Operators just pay for the data that is sent to Confluent Cloud and the system automatically scales based on the throughput and desired retention period. No more pre-provisioning of compute and storage, no more waiting to scale.

For Confluent Platform, Tiered Storage allows Kafka to recognize two tiers of storage: local disks and cost-efficient object stores, such as Amazon S3 or Google Cloud Storage (GCS). This enables you to offload older topic data from expensive broker storage to inexpensive object storage, significantly reducing storage costs. This is useful for use cases involving historical data, such as storing financial transaction data for compliance reasons.

Self-Balancing Clusters is another feature which fully automates partition reassignments, ensuring your data is evenly distributed across a cluster to optimize performance and resource utilization. Finally, Multi-Region Clusters reduce the infrastructure needed for multi-datacenter Kafka deployments by leveraging Kafka's internal replication engine to durably replicate partitions across multiple datacenters with a single stretch cluster, rather than needing to deploy a separate Kafka Connect based replication tool.

Efficient Operations at Scale

Operating Kafka in production also requires a significant and ongoing investment of engineering resources. Responsibilities for DevOps and platform teams include configuring, deploying, securing, upgrading, scaling, and maintaining the uptime of the cluster. These tasks are manual, complex, and time-intensive—time that could otherwise be spent on high-value projects for the business. They also often require building automation tools to simplify management and ensure the platform runs reliably, but developing those components in-house comes with risks and requires additional investment of engineering resources.



"With Confluent, we've been able to optimize our in-house Kafka-support resources by ~50% and redeploy our engineering & development talent to creating innovative and value-added features for our customers instead. Additionally, engaging Confluent gave us the peace of mind that comes with minimizing downtime."

— JOSH LITTLESUN, DIRECTOR OF TECHNOLOGY, ROADSIDE

Confluent Cloud allows you to completely offload cluster management to the world's foremost experts with over one million hours of Kafka experience, ensuring your best engineers stay out of low-level infrastructure management and instead focus on up-the-stack development that delivers competitive advantages for your business.

With Confluent Cloud, you can:

- Quickly deploy your Kafka cluster within minutes and elastically scale to meet all of your data streaming workloads
- Effortlessly connect your critical data sources and sinks to build streaming data pipelines with fully managed connectors that eliminate developer time and ops burden
- Easily build high-quality, reusable data streams with the industry's only cloud-native, serverless Flink service

Confluent Platform can also help you reduce the time and resources spent both building out components to make Kafka an enterprise-ready data streaming platform and managing the platform day-to-day.

Designing, building, testing, and maintaining Kafka connectors can take up to ~6 engineering months each and connecting to proprietary systems can require high licensing costs. Surfacing important cluster health metrics also requires identifying which metrics need to be highlighted to effectively monitor and manage the cluster. That's why Confluent Platform includes 120+ pre-built and battle-tested connectors and Health+ for cloud-based alerting and monitoring, so developers can focus on what matters most. Moreover, Confluent for Kubernetes simplifies running Confluent Platform as a cloud-native system on Kubernetes to further reduce operational burden by automating configuration, deployment, upgrades, and scaling.

Confluent Support, Professional Services, and Training underpin our offerings to further reduce the time needed for both platform development and operations. With over 1 million hours of technical experience with Kafka, Confluent has helped customers quickly ramp up on Kafka fundamentals, integrate Kafka into their existing architecture, ensure the platform's configuration meets best practices, and greatly reduce the time spent on troubleshooting and maintenance.

Minimized Risk and Downtime

The impact of risks, such as downtime, performance degradation, and security breaches, should not be understated. According to a [Forrester study](#), planned downtime costs enterprises \$1.5M per quarter. Unplanned downtime costs 35% more per minute than planned

downtime. And a recent study by IBM found that the average data breach in 2022 costs an organization \$4.35 million, across detection and resolution costs, fines and reparations, and reputational harm to the business.

Some organizations may claim little downtime in their self-managed environments—but even they suffer significant performance issues, which can lead to prolonged periods of lost productivity and reduced customer satisfaction.

As an example, we worked with an organization with 2,000 total employees, 25% of which were software engineers—25% of these developers worked on Kafka-related projects. This means that ~125 employees were impacted whenever Kafka experienced performance-related issues.

When we calculated the fully loaded cost of these resources and the impact to productivity, we were able to calculate the total cost was around \$230K per incident, not to mention the ongoing frustration and distraction of such issues for your valuable engineering resources.

It takes significant investment and expertise to maintain a reliable platform with high availability—factors that only grow over time. Confluent Cloud alleviates these burdens with a 99.99% uptime SLA, along with no disruptions for upgrades and planned maintenance. And Multi-Region Clusters for Confluent Platform allows operators to increase data durability with near-zero RPO and RTO, all while automating client failover in the event of a disaster. In addition, Confluent's 24x7 committer-led support helps prevent issues and enables fast recovery.

Finally, when it comes to security breaches, the safety of your data is obviously paramount. It can be costly and risky to manage it all by yourself with open source Kafka. Confluent comes with enterprise-grade security, compliance, and governance controls to ensure your data is always safe and secure.



"In terms of hard costs, Confluent reduced our bottom line from the infrastructure cost perspective, reducing licensing and physical hardware. We reduced our projected yearly cost by 69%. And the months of planning we would have to put into meetings and projections are gone, so our management oversight is reduced from a soft-cost perspective."

— JUSTIN DEMPSEY, SENIOR MANAGER, SOFTWARE DEVELOPMENT, CLOUD & INFORMATION SERVICES

Accelerate Time-to-Value from Months to Minutes

Beyond reducing your Kafka TCO, Confluent can also rapidly accelerate time-to-value for both new and existing data streaming use cases. While Kafka is the foundation for the modern data stack, the open source project doesn't offer everything you need out-of-the-box. Here are just a few of the common activities needed to implement data in motion use cases end to end:

1. Kafka brokers and ZooKeeper need to be manually configured and deployed
2. Connectors need to be developed, tested, and maintained to integrate your critical apps & data systems
3. GUI tools need to be integrated and relevant metrics need to be surfaced to monitor and alert on cluster health
4. Security features for access control, encryption, and auditing need to be custom-built
5. Disaster recovery tools and failover logic have to be created to protect your business against downtime and data loss

Traditional Kafka Development



Only after building all of the foundational components for Kafka can you deploy truly production-safe streaming applications and pipelines. Many businesses will also need to hire engineers familiar with Kafka to start this process, which adds to the timeline. In sum, this whole process can take up to two years to finally get to production at scale, significantly slowing time-to-value and eliminating any first-mover advantage in the market.

Development with Confluent



To help companies realize the full value of data streaming, Confluent offers an elastically scalable and globally available service to accelerate this process dramatically. Rather than spend valuable development cycles on building foundational tooling for Kafka, customers can leverage our complete platform to get to production quickly, securely, and reliably.

As a fully managed, cloud-native service that's deployable in minutes and available across all major public clouds, Confluent Cloud allows you to get started for free and keep costs low with pay-as-you-go pricing that scales as you grow. With Confluent Cloud, we do all the work for you—clusters are provisioned instantly and maintenance is seamlessly managed. You can start streaming data on the day you sign up—in minutes instead of weeks or months, and without the operational overhead.

As a self-managed option, Confluent Platform includes Infrastructure-as-Code tooling to programmatically deploy clusters, 120+ pre-built connectors to popular sources and sinks, management and monitoring tools specifically designed for Kafka, enterprise-grade security and disaster recovery features, and more.

Finally, Confluent's committer-driven expertise also accelerates time-to-value with Kafka. For users less experienced with Kafka, initial flaws in the configuration and architectural design of clusters can lead to increased troubleshooting, tech debt, and delayed project timelines. Confluent provides a suite of support and services to implement optimal design decisions to deploy clusters that are scalable, secure, and resilient. Additionally, our education offerings can quickly build internal Kafka expertise throughout an organization so more developers can start building real-time applications that leverage data in motion.

Maximize your ROI with a 3X Return on your Kafka Spend

A ROI exercise models a financial return on an investment—and may include forecasting the benefit, or return, of implementing a data streaming use case. Any credible ROI calculation should include a range of outcomes, including lower/upper bounds (sometimes referred to as a Monte Carlo simulation). We always try to model a range with a conservative estimate in the middle. This is the best approach for forecasting potential events despite assumptions and factors outside of real-world control.

A Confluent subscription can pay for itself more than 3X over, driving savings across infra costs, FTEs needed, and set up and operating costs.

	Confluent Cloud	Confluent Platform	Max Change from Kafka
Reduction in Infrastructure	\$1.85M	\$0.80M	100%
Reduction in Development	\$0.30M	\$0.25M	(Cost Savings–Subscription)/ Subscription
Reduction in Operations	\$0.23M	\$0.15M	
Reduction in Support and Services	\$0.14M	\$0.14M	100%
Subscription Cost	–\$0.37M	–\$0.47M	N/A
ROI on Subscription	5.9x	3.7x	

ROI is even greater when accounting for harder to quantify costs and benefits, such as reduced downtime, security risk, and performance degradation, along with an accelerated time-to-market.

Additional benefits that may not be fully captured in a TCO calculation include increased developer velocity, improved business agility, reduced overall risk, and perhaps most importantly, freeing key resources from the operational burden of setting up and operating Kafka clusters. Eliminating the deployment and operations of Kafka frees valuable talent to focus on the challenges and problems that are core to your business.

Our unparalleled experience with helping over 4,000 customers (and counting) realize the value of Data in Motion across industries and maturities means that we can help you unlock the full value of your data streaming use cases from inception to large scale. This includes future-proofing the business and modernizing the tech stack enabling a microservices-based architecture—for improved integration across other parts of the business that might later adopt the platform. This also means showing early success to build momentum and drive sponsorship, building a roadmap of subsequent real-time use cases to pursue, and driving broader understanding and adoption of the technology.

Confluent Cloud versus MSK

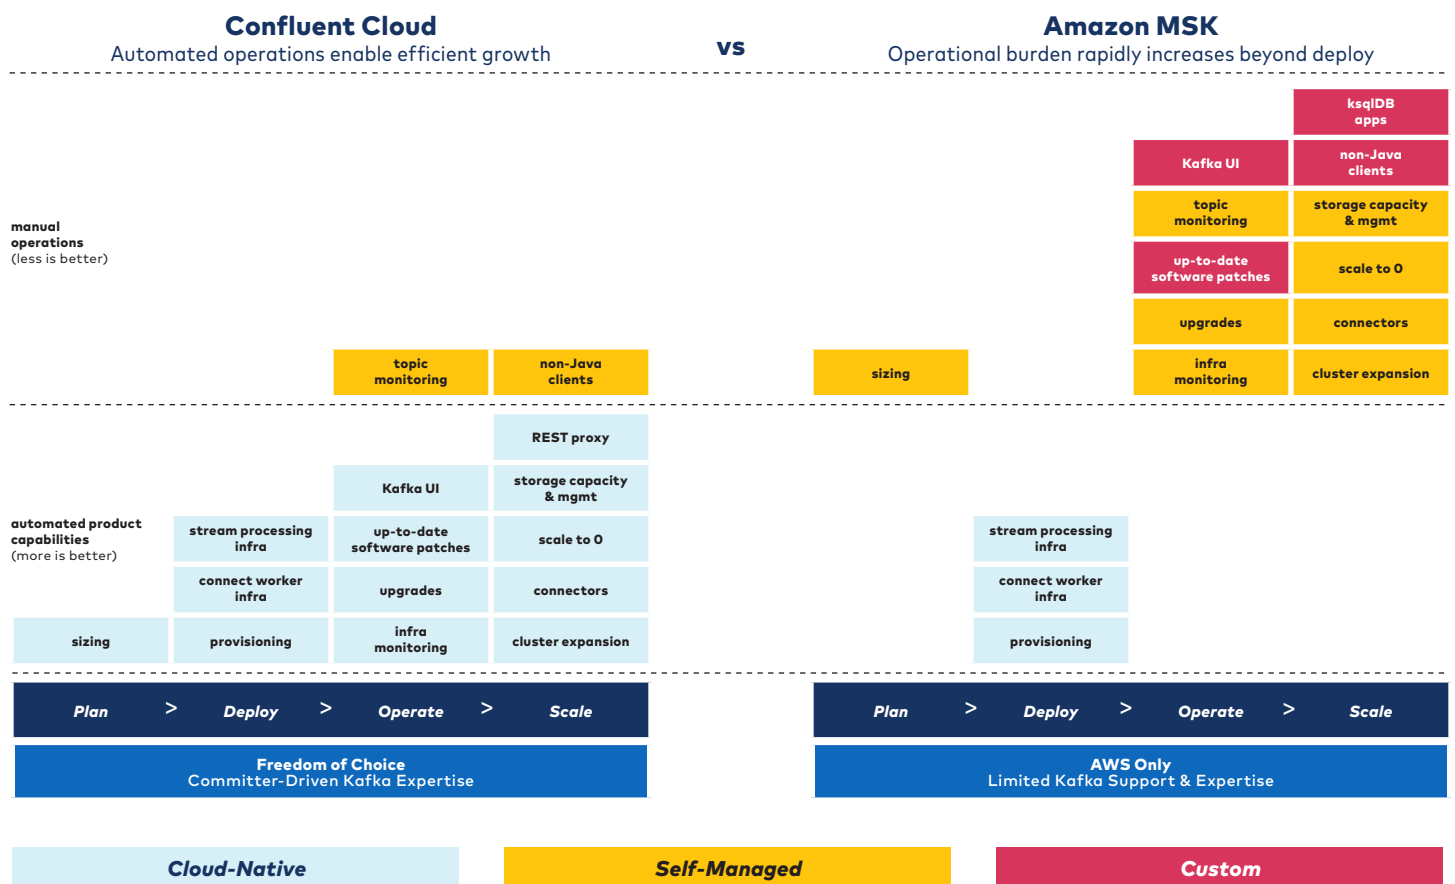
Many Kafka services will call themselves fully managed—but the definition of fully managed can mean different things to different services.

On one end of the spectrum you have a truly cloud-native experience where product features and automation make operational tasks effortless, enabling teams to focus on their applications. This is the case with Confluent Cloud, which offers a true cloud-native experience across your Kafka journey all the way from planning to scaling. On the other end, you have services with capabilities that are not available or supported by the vendor that lead to teams having to build custom code without any support.

And in the middle you are going to have what services will try to sell you as fully managed when really they require manual operations with additional tools. Many managed services promise to alleviate the operational burden of Kafka, but in reality most of them—including Amazon MSK—are only able to meet your needs early in your data streaming journey, with easy to deploy capabilities. However, these lack the operational automation capabilities that truly enable teams to step away from managing Kafka day-to-day and focus on their applications. Although Amazon introduced a serverless offering as well (MSK Serverless) that addresses some operational gaps, it lacks many mission-critical components, such as limited connectors, stream processing, Kafka-specific security & governance, and more.

Ideally, every step across your event streaming journey is cloud-native in nature so your team does not get overwhelmed with Kafka operations that should be automated by your managed service. So, when we compare options we tend to get more technical as we're comparing specific solutions—and we can split this across the project lifecycle of: PLAN > DEPLOY > OPERATE > SCALE.

The diagram below compares Confluent Cloud across this lifecycle with Amazon MSK, illustrating the differences between a cloud-native and cloud-hosted solution.



Because MSK is not fully-managed and cloud-native, the costs and operational burden increase especially when you go beyond the deploy phase of the data streaming project lifecycle. What follows is the impact on the cost categories and comparison with Confluent Cloud.

Infrastructure

Because MSK uses node-based pricing, you always pay for provisioned infrastructure, both compute and storage, for every cluster, even for those that may have low throughput during development. You also have to allocate time and technical resources to run multiple performance tests to pick broker types and broker count. MSK Serverless offers throughput-based sizing supporting up to 200MBps ingress and 400MBps egress.

In terms of provisioning, both Confluent Cloud and Amazon MSK (both provisioned and serverless) can help gain access to quick, on-demand provisioning as part of a cloud-native experience. What's unique about Confluent Cloud is that you pay for usage versus infrastructure so that your team can optimize costs throughout the Kafka journey.

Confluent Cloud also uses a throughput-based model as opposed to the node-based model to abstract the complexity of sizing the underlying Kafka infrastructure. With a throughput-based model, you reduce costs from upfront oversizing and only pay for storage used.

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."

— MATT CODDINGTON, SENIOR DIRECTOR OF DEVOPS ENGINEERING

Operational Costs

With a cloud-hosted service such as Amazon MSK you have to consume low level metrics with the goal of self monitoring, diagnosing and correcting issues related to cluster performance.

MSK also enables customers to scale using their console or an API, but that process requires active manual monitoring of your environment. MSK also requires a manual process to rebalance and reassign partitions, which is a non-trivial task while scaling out. Also, MSK does not allow you to scale in. MSK Serverless does offer elastic scaling with no need to rebalance the cluster. However, there is no solution to upgrade when exceeding the given quota.

With Confluent Cloud, there's proactive monitoring and reduction of monitoring costs with no infrastructure level metrics to export and analyze. For topic monitoring, Confluent Cloud offers free, pre-aggregated metrics while MSK is pay to consume and requires manual aggregation. And in terms of scaling, Confluent Cloud offers elastic scaling that reduces operational burden and costs with scale-to-zero pricing on Basic clusters.

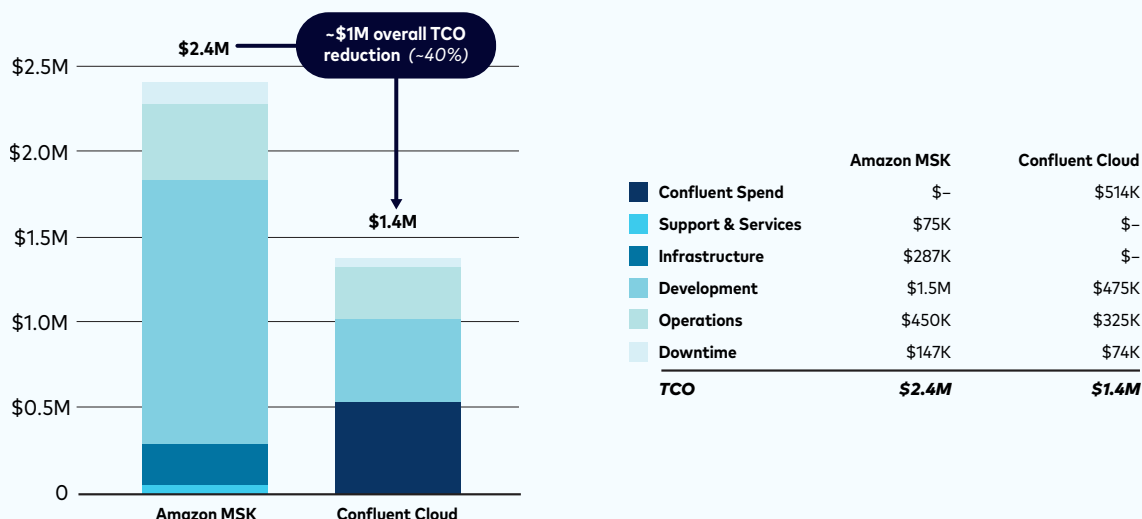
Intangible Costs

There are several factors which increase the risk of downtime when it comes to MSK. First, because it's possible to underestimate storage of the broker up front, it's possible that a spike in demand leads to an outage and associated downtime costs. Next, manual data balancing increases the risk of downtime due to potential error-prone manual scaling operations.

And a lack of rolling upgrades and up-to-date patches, limited version support, and SLA exclusions for Kafka failures means that critical issues could remain unresolved for extended periods of time.

On the other hand, with Confluent Cloud you can seamlessly scale your cluster to meet your streaming needs so there is no risk of underestimating storage. Automatic data balancing means that all Confluent Cloud Kafka clusters are continually optimizing data placement to balance load—no error-prone manual operations that could result in downtime or performance degradation required. Finally, rolling upgrades to the latest stable Kafka version that includes strategic patches ahead of scheduled Apache releases and a 99.99% uptime SLA de-risks your investment and prevents business disruption.

In an example assessment, when analyzing a three-year cost difference, we modeled \$1M cumulative savings with Confluent Cloud versus Amazon MSK which is a 40% reduction in data streaming TCO.



Conclusion

Regardless of your familiarity and maturity with Kafka or data streaming, Confluent can drive cost savings that cover the cost of a subscription several times over. When factoring in the accelerated time-to-value with new use cases, the ROI on choosing Confluent over the open source technology or other cloud-hosted Kafka services is clear and significant. Of course, this is in addition to ensuring the success of your data in motion initiatives and maximizing the value that they deliver to your business.

If you would like a customized TCO assessment tailored for your specific Kafka use case, please [contact Confluent](#) to start the discussion. And for those who would like to try our complete, cloud-native platform, visit our [Get Started](#) page today to try Confluent Cloud for free.

About Confluent

Confluent is pioneering a fundamentally new category of data infrastructure focused on data in motion. Confluent's cloud-native offering is the foundational platform for data in motion—designed to be the intelligent connective tissue enabling real-time data, from multiple sources, to constantly stream across the organization. With Confluent, organizations can meet the new business imperative of delivering rich, digital front-end customer experiences and transitioning to sophisticated, real-time, software-driven backend operations. To learn more, please visit www.confluent.io.