

GOVERNMENT

Driving Data to DoD Software Factories with Mission-Ready Apache Kafka[®]

The Challenge

Threats from our adversaries are constantly evolving; the US Department of Defense (DoD) must also evolve at the speed of relevance. As an example, the US Air Force has embraced DevSecOps, Continuous Integration (CI), Continuous Delivery (CD), and Continuous Authority to Operate (CATO) as part of an effort to boost agility and responsiveness to changing requirements and speed of operations. This initiative has become a priority across the PlatformONE-affiliated software factories and increasingly across the entire DoD.

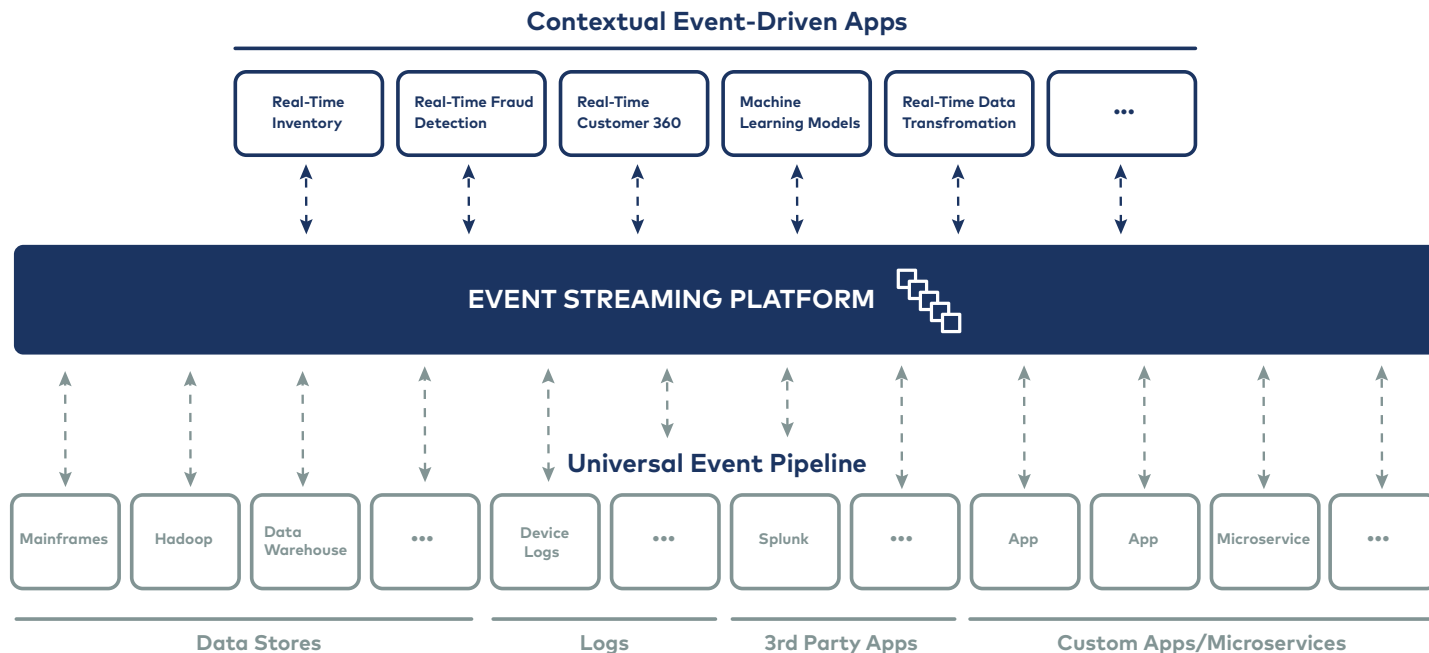
In the move towards agile DevSecOps, older waterfall development methodologies must be overhauled and improved. The DoD recognizes the advantage of being able to fail fast and rapidly iterate on new mission capabilities in order to best support the war fighter. This requires near-immediate access to data services for rapid app development and the ability to stream data to apps in real-time so that users can respond as events occur. Meeting this demand means having the ability to easily create, use and repackage data services across all data producers and consumers. The resulting modular approach to data improves accountability, program effectiveness, and ultimately mission success.

| | |
|--------------------------------|----------|
| The Challenge | 1 |
| The Solution | 2 |
| Flexibility. | 3 |
| Durability | 4 |
| Security | 5 |
| Summary | 5 |

The Solution

Confluent Platform is the mission-ready packaging of Apache Kafka, which has emerged as the open-source industry standard solution to this problem. Kafka is a key technology cited by the DoD DevSecOps Initiative as part of its Product Stack. It is an event-driven data streaming and messaging architecture that integrates applications while removing dependencies between producers and consumers of data. It facilitates the asynchronous data exchange between processes, applications and servers, and can process trillions of events per day. Kafka is currently used by over 100,000 organizations globally.


Confluent Platform completes the event streaming platform and adds the flexibility, durability, and security required for complex, large scale mission operations. Confluent is a high-speed, cloud-native, secured data back-plane integrating container-run microservices in CI/CD pipelines. Not only do commercial companies such as Walmart, Square, Expedia, and Unity rely on Confluent to run their business, the DoD does too. Confluent Platform has been through the STIGs process and granted ATO for many production deployments on classified networks, across security guards and in DIL network environments.



Confluent was founded by the original creators of Kafka in response to the explosive growth and broad adoption of Kafka. Confluent has over 1 Million hours of experience in support and development of the Kafka platform and executes a large majority over 80% of the code commits to the Apache Kafka baseline. This expertise, combined with the features Confluent has built to extend the Apache Kafka core, consistently reduces risk to program cost, schedule and performance, minimizing the time lost to troubleshooting complexities of mission architectures.

Flexibility

Flexibility in a data integration scenario means being able to rapidly incorporate emerging data sources, formats and velocities. It also means being able to scale up in a predictable and linear fashion in the face of explosive growth in data throughput-rates. It means being usable by as broad a range of skill sets and user types as possible. Lastly, flexibility in integration means being able to quickly and easily integrate data between a large number of data sources and destinations with low or no-code required, combining data streams as dictated by the requirement.

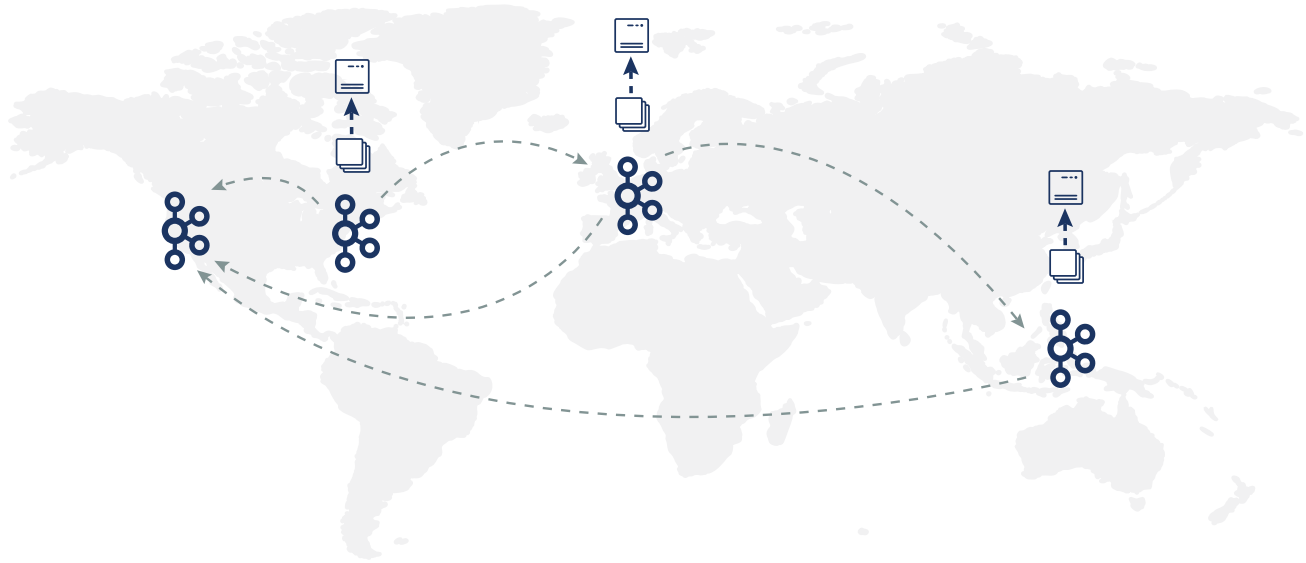


Confluent Platform extends the Apache Kafka Connect ecosystem with more than 100 supported and Confluent-verified connectors already implemented. Using them to integrate to your RDBMS, data lake, HDFS, message bus, or cloud store is as easy as configuring a few parameters. Talented Java developers are in short supply, so Confluent has expanded the Confluent Platform to support Python, Go, .Net and C++ as development languages. Further, the KSQL feature allows analysts to create new data streams with a scripted language that works like SQL.

Confluent allows its customers to use an iterative, milestone-focused approach where projects start modestly and scale later to meet full mission requirements. That's because Kafka is a cloud native event streaming platform that scales by simply adding new brokers to the cluster as required. Flexibility also means control over deployment. Confluent deploys on simple bare metal, on Virtual Machines, or in Kubernetes pods automated by Confluent's Operator. Deployment of software from configuration managed, repeatable code-driven configuration is a best practice of the DevSecOps agile and iterative methodology; Confluent has embraced it with these supported tools and feature capabilities.

Durability

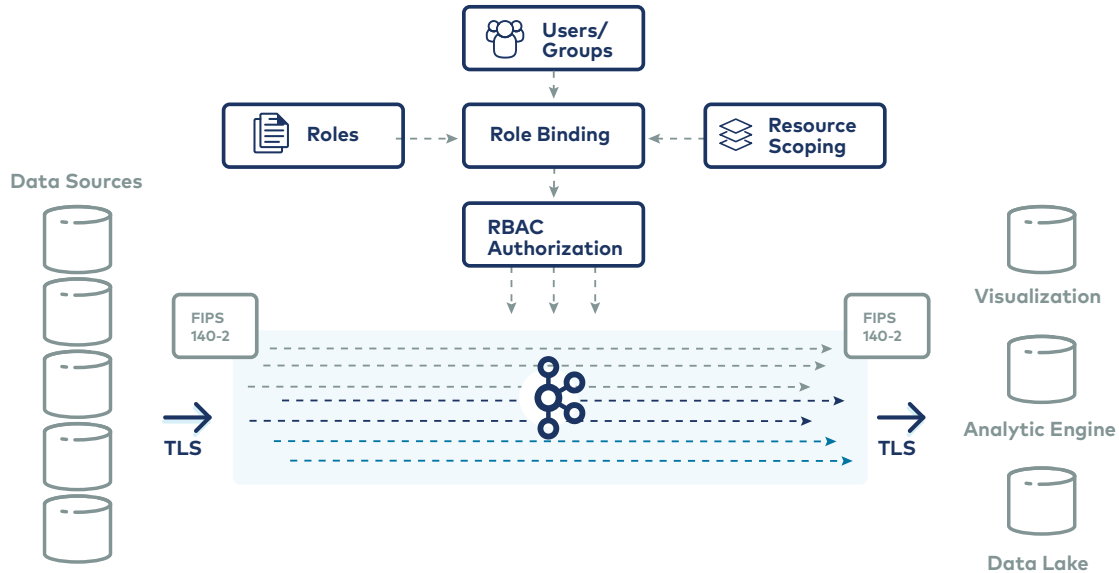
Messaging has long been used to integrate systems, especially for real-time use cases, but there have been hard limitations on its applicability because of the ephemeral nature of the messages. Once processed, the data was lost, leaving a significant resilience gap and no ability to replay the data. Kafka solves this problem by persisting messages in an immutable continuous commit log, for a configurable length of time, providing auditability and accurate, ordered replay of data to consumers. Additionally, Kafka supports atomic transactions and exactly-once message semantics to ensure that even in partial-failure conditions delivery of messages can be guaranteed without delay or duplication.



In addition to preserving real-time data, durability in enterprise also means resiliency in the face of outages of servers or even whole data centers. Kafka is a horizontally scalable, distributed and resilient platform designed to withstand outages and automatically rebalance and redistribute workloads to ensure uptime and maximum performance—even in a degraded infrastructure or network state. Confluent Platform further extends this with Replicator, enabling Kafka clusters around the world to stay in sync—even with DIL (Disconnected, Intermittent, Limited) network environments. Mission functions are thus able to operate at near-zero RPO/RTO. Confluent also adds the ability to create Multi-Region clusters, stretching Kafka over thousands of miles. This creates a single logical application for users and applications who seamlessly use the fastest-available copy of data, continuing to operate accurately if one of the sites goes down. With these innovations, Confluent Platform gives Kafka the resiliency to meet global and real-world mission challenges.

Security

Confluent brings the most critical set of enterprise differentiators in the form of security around Kafka. Confluent deployments have been through STIGs process with multiple ATOs on various DoD networks.



Confluent Platform adds security features such as FIPS 140-2 compliant mode, Secrets to encrypt authentication traffic, integration with identity management tools like Active Directory and LDAP, and Role-Based Access Control to restrict access to data in topics by user-role. Confluent’s Audit Logging makes it possible to track in real time which data has been accessed, which configurations were changed and by whom. These Confluent Platform capabilities elevate the open-source Kafka capabilities to address critical DoD security requirements.

Summary

As the DoD continues to invest in DevSecOps as a culture and approach to rapidly meeting the warfighter’s needs, it needs secure yet widely available access to cloud native infrastructure. These tools enable microservice architectures using containers on dynamically provisioned infrastructure with Continuous Integration and Delivery. Kafka is the open-source industry standard bringing these capabilities to the stack. Confluent Platform delivers Mission-Ready Kafka, making it more secure, more consumable and more resilient. It extends and completes the Apache Kafka event streaming platform with advanced Flexibility, Durability, and Security features that accelerate application development and integration, enable event transformations through stream processing and simplify enterprise operations at scale, while meeting the strict security guidelines required by our DoD customers.

Nobody and no company knows Kafka as well as Confluent does. We provide service and training options to quickly enable development and administration teams to exploit the capabilities Confluent Platform brings to the mission. Learn More about Confluent at www.confluent.io/resources/ and discover how event streaming will unlock real-time mission-focused applications.

To discuss how Confluent can enable your mission, visit confluent.io/government.

Ready to get started? Contact a Confluent expert today

Email us at publicsector@confluent.io | Or visit www.confluent.io/government/ for more details