

# Build an Integrated Trading Ecosystem with Data in Motion

**Brian Pyle** Managing Director, FSI Group

**Chris Matta** Field CTO and Principal Solutions Engineer

**Jeff Ferguson** Senior Product Marketing Manager

- Data in motion use cases for financial services** . . . . . 2
- Background: The trading stack problem** . . . . . 3
- Overview: The trading stack solution** . . . . . 3
- Connectors: Easily sync systems and applications** . . . . . 4
- Stream processing = data movement + intelligence** . . . . . 4
- Hybrid and multi-cloud: A single pane of glass for your entire global business** . . . . . 7
- Trading use cases: What your peers are doing** . . . . . 9
- Set your data in motion with Confluent** . . . . . 11

# Data in motion use cases for financial services

Data in motion is data's natural state.

All data originates as continuous streams of events that represent things happening in the real world. Motion more accurately describes the world than rest. Yet, most people tend to think about data as a static object, stored away somewhere to be used for some purpose at a future point in time. This shift in perspective about data as something in motion instead of something at rest empowers us to reimagine our products and services as intelligent, contextual, data-driven systems that continuously react to things as they happen. Event-driven systems result in experiences that delight customers, provide operational efficiency gains, reduce costs, and increase top-line revenue by fueling new capabilities. Financial services companies that merely use software are being displaced by fintech startups that fundamentally are software. Building better software is not a matter of luxury. This is a matter of life and death for businesses.

All 10 of the top 10 U.S. banks are Confluent customers. These banks choose to partner with Confluent because our enterprise platform for data in motion allows them to access and utilize all of their data in real time to build software applications that meaningfully improve the way they do business. The image below shows the most common data-in-motion use cases we see from our customers in the financial services vertical.



Some of these use cases are unique to financial services, while others are applicable across industries. This paper showcases the value of data in motion with a focus on trading and provides:

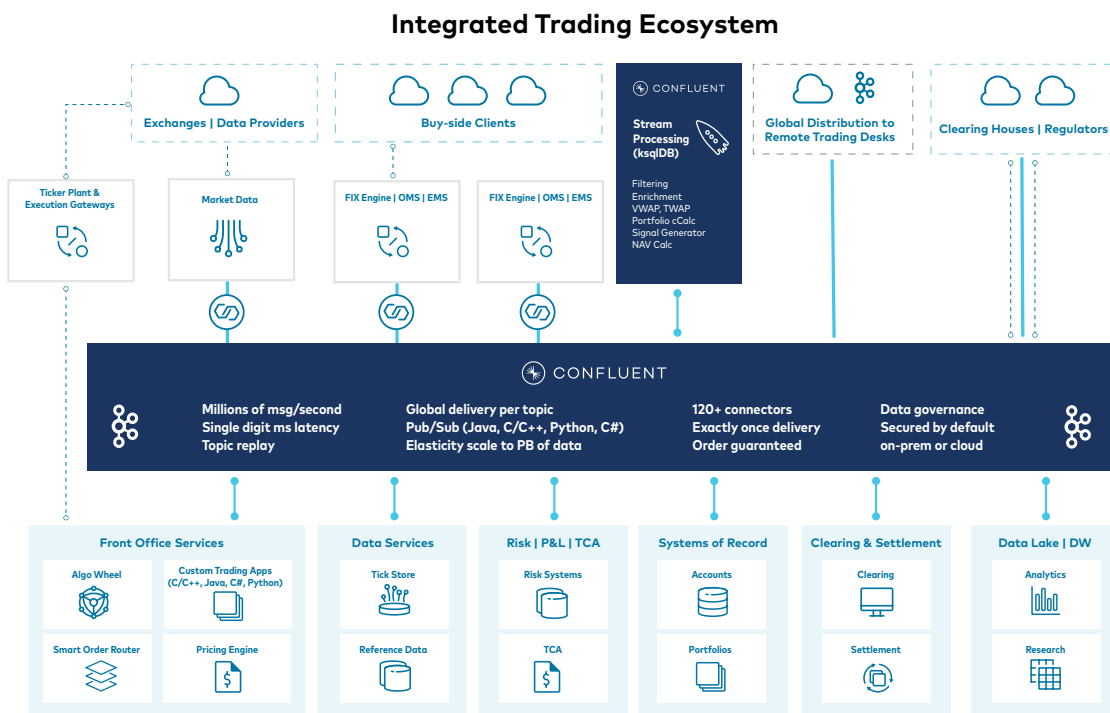
- Background on the trading ecosystem
- An overview of how a data-in-motion platform unifies siloed systems and applications in the trading ecosystem
- Detail on how stream processing powers intelligent automation in trading
- Insight into how hybrid and multi-cloud architectures can unify computing environments across on-premises, cloud, and global regions while maintaining consistency
- Confluent customer trading use cases and benefits

# Background: The trading stack problem

Sell-side firms' technology stacks have historically been built to support specific asset classes. Fixed-income trading desks, equities, and derivatives each use a different set of systems and applications to support their operations. This organic evolution of disparate sets of systems, each for a different purpose resulted in "asset class islands," which have grown over the years as firms' trading desks expanded globally and branched into new asset classes. Data, application, and system silos have resulted in pressure on trade support staff, operations, risk, compliance, and IT teams to keep pace with demands from the front office and an ever-changing set of regulations. These pressures have resulted in higher settlement failure risk, inability to respond to client reporting requests, and an unacceptable number of manual exceptions for trade processing. An increase in trading volumes in 2020, while fueling trading profits, has also exposed areas of opportunity for technology modernization.

# Overview: The trading stack solution

The diagram below illustrates a next-generation trading platform that has unified siloed datastores, systems, and applications. This singular, unified architecture empowers users to harness and utilize data in motion, wherever that data lives. It's worth noting here that the word "users" does not refer only to human beings, but also to systems and applications interested in consuming infinite streams of data for any number of purposes. Data in motion is the lifeblood of real-time trading applications. It is the key to building simpler, faster, more cost-effective software systems that reduce operational costs and increase top-line revenue. By enabling new services and delighting customers with better experiences, organizations that unify their data at rest and their data in motion to power innovative software applications become more competitive in their respective markets and disrupt their competition instead of being disrupted by them.



In this diagram, market data from exchanges, data providers, and low-latency ticker plants is captured using Confluent as a centralized platform for data in motion and distributed to stream processing engines, remote trading desks, clearinghouses, front office services, and various downstream datastores. The unification of these front-office and back-office systems is possible within and across asset classes, integrating these "asset class islands" globally and without disruption to the business. Confluent is capable of handling millions of messages per second with single-digit latency. And perhaps most importantly, point-to-point communications are eliminated with a centralized platform for data in motion.

But haven't we seen this vision before? Haven't banks and other financial institutions invested in technologies like messaging systems (TIBCO EMS, IBM MQ), ESBs, and EAI technologies? Couldn't we just replace the Confluent image above with another technology and vendor's name in the middle? While there are some parallels to other messaging, pub/sub, integration, and stream processing technologies, there are some important architectural differences.

Confluent is very different from traditional message queues where messages are ephemeral and deleted after consumption, or from "fire and forget" multicast style systems. Unlike traditional messaging systems, Confluent persists data to disk, immutably and in order, for any period of time. This enables replay, ordering guarantees, exactly once semantics, and the ability to use Confluent as a datastore. While Confluent is not designed to replace purpose-built databases, data stored in Confluent does alleviate pressure from external systems and enable applications to utilize real-time data in Confluent instead of needing to query external systems.

In addition, legacy messaging systems and the technologies built on top of them (ESBs) are fragile and don't scale to meet the volume and velocity requirements of modern trading firms. Confluent, in contrast, is a horizontally scalable and fault-tolerant platform, making it highly performant and resilient. Confluent is used in production by the world's largest banks and scales to the needs of the world's most data-intensive users across financial services, tech, retail, gaming, telco, government, and other verticals that demand scalability, performance, and the kind of rock-solid stability that can support critical business systems and operations.

## Connectors: Easily sync systems and applications


Data pipelines are easily created using [connectors](#), which provide seamless integrations to popular data sources and destinations. Data is produced and consumed in a publish and subscribe fashion using a variety of languages including Java, C, C++, C#, and Python. Compatibility and governance of schemas are maintained using Confluent's [Schema Registry](#). Data can be written to Confluent and applications can consume data from Confluent without the upstream and downstream producers and consumers needing to know anything about each other.

Order flow coming from FIX engines, execution and order management systems is captured using Confluent's JMS connector for TIBCO EMS, Solace, MQ, and other legacy messaging systems. This data flows continuously through Confluent for distribution to additional systems for risk, clearing and settlement, reporting (e.g., OATS/CAT) and to data lakes and data warehouses for analytics.

With over 100 pre-built connectors for the most popular data sources and sinks, the connect framework is a key part of Confluent's data-in-motion platform and radically simplifies the integration of systems and applications, setting data in motion across an organization.

## Stream processing = data movement + intelligence

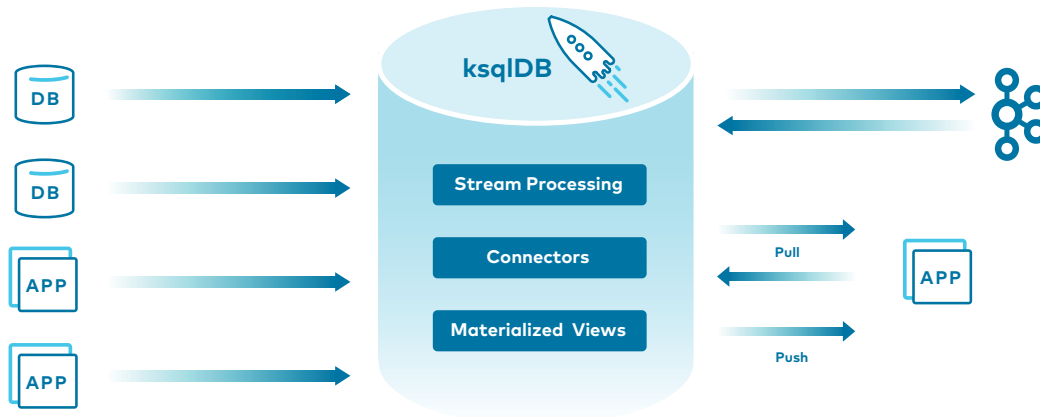
Once an organization implements a data-in-motion platform, it starts to have a natural network effect. The platform provides more value to the business by connecting to more systems and capturing more use cases and, of course, more data. Confluent starts to function as the central nervous system for an organization. And if the unification of systems and seamless flow of data across an organization is the central nervous system, then stream processing is the brain.



As data in motion flows through Confluent, stream processing adds intelligence using [ksqlDB](#), a SQL-based engine with persistent integrated storage. ksqlDB provides a familiar and lightweight SQL syntax that developers can use to construct stateful materialized views over data in motion. ksqlDB further empowers its users to filter, enrich, transform, and join data in motion on real-time and historical data, all with an elastic and fault-tolerant approach. Streaming joins allow multiple event streams to be joined to model logic or provide dimensions for filtering. Perhaps a use case requires enriching trade data with reference data, like adding a CUSIP or ISIN code to each trade message. ksqlDB can easily accomplish this with a few lines.

```
CREATE STREAM ENRICHED_TRADES AS
SELECT
  A.symbol as symbol,
  B.cusip as cusip,
  B.isin as isin,
  B.description as description,
  B.sector as sector,
  A.bid as bid,
  A.ask as ask,
  A.size as size,
  A.lastTradeTime as lastTradeTime
FROM MARKETDATA A JOIN MARKETDATA_REFERENCE B ON A.symbol = B.symbol
```

Any ksqlDB query results in a stream of events that can be pushed to clients or subscribed to as a topic. Confluent calls these “push queries.” Instead of polling for an answer like a traditional datastore, ksqlDB allows clients to subscribe to a query and have records pushed to them. For instance, maybe a portfolio manager is only interested in the securities in their portfolio. The ksqlDB service can also materialize stateful data in [tables](#) for aggregations that support traditional in-memory pull queries. Combining the data capture, processing, and presentation layer to a single service instead of various processors, caches, and datastores drastically simplifies the overall architecture. Data capture, streaming ETL, persistence, and the query layer are all handled by ksqlDB.



Using SQL for stream processing accelerates developer velocity. It frees up valuable engineering resources to focus on building things that are core to an organization’s business instead of writing and managing large volumes of complex application code associated with low-level data infrastructure. ksqlDB users can write push queries that take action when something interesting happens, like a trading volume increase greater than a certain percentage. Users can also write pull-based queries to look up a real-time snapshot of something, like the current valuation of a portfolio.

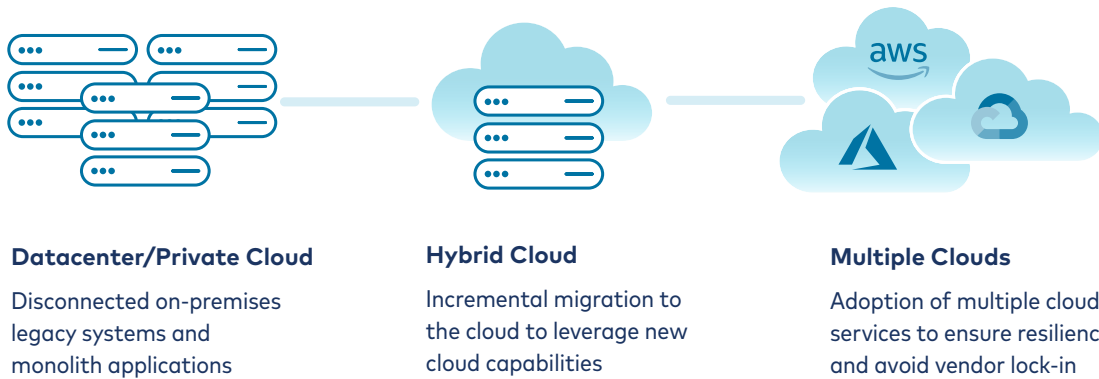
For example, let’s say a trading desk wants to calculate intraday volume-weighted average price. We can use ksqlDB to analyze data in motion that is *already* flowing through Confluent to calculate VWAP. The example below shows how simple this is to accomplish using ksqlDB:

```
CREATE TABLE TRADES_VWAP
AS SELECT
  TRADES_STREAM.SYMBOL SYMBOL,
  ((SUM(TRADES_STREAM.PRICE) * SUM(TRADES_STREAM.VOLUME)) / SUM(TRADES_STREAM.VOLUME)) as VWAP
FROM TRADES_STREAM TRADES_STREAM
WINDOW TUMBLING ( SIZE 1 MINUTE )
GROUP BY TRADES_STREAM.SYMBOL
EMIT CHANGES;
```

Consider what would be required to create a real-time stream processing application that could join event streams, utilize time windowing, and leverage a push and pull query layer. In the example above, a developer was able to analyze and act on data in motion in minutes using a few simple SQL statements. The simplicity of this approach to stream processing results in faster, less expensive, and lower risk operations that a wider range of users can leverage to build real-time applications on top of data in motion to do things like calculate intraday VWAP and then take action when some set of criteria is met. Like Albert Einstein once said, “Everything should be made as simple as possible, but no simpler.”

# Hybrid and multi-cloud: A single pane of glass for your entire global business

The data-in-motion platform described above is great for unifying systems and applications in a single environment, but what about scenarios with multiple environments? Many organizations have one or more datacenters on-prem, as well as one or more cloud environments. Fortunately, the “single pane of glass” for data infrastructure described above is not limited to a single computing environment. It can also unify different environments themselves.



Most organizations are embracing a cloud-first strategy, but cloud migration is a long-term journey. In some cases, on-prem applications are inextricably tied to legacy systems, making cloud migration difficult or impossible. In these scenarios, the most common solution we see among enterprises using Confluent is to bridge their on-prem deployments to the public cloud, creating a hybrid architecture. This enables the utilization of data from on-prem systems in cloud-native, real-time applications and progresses enterprises on their cloud journey.

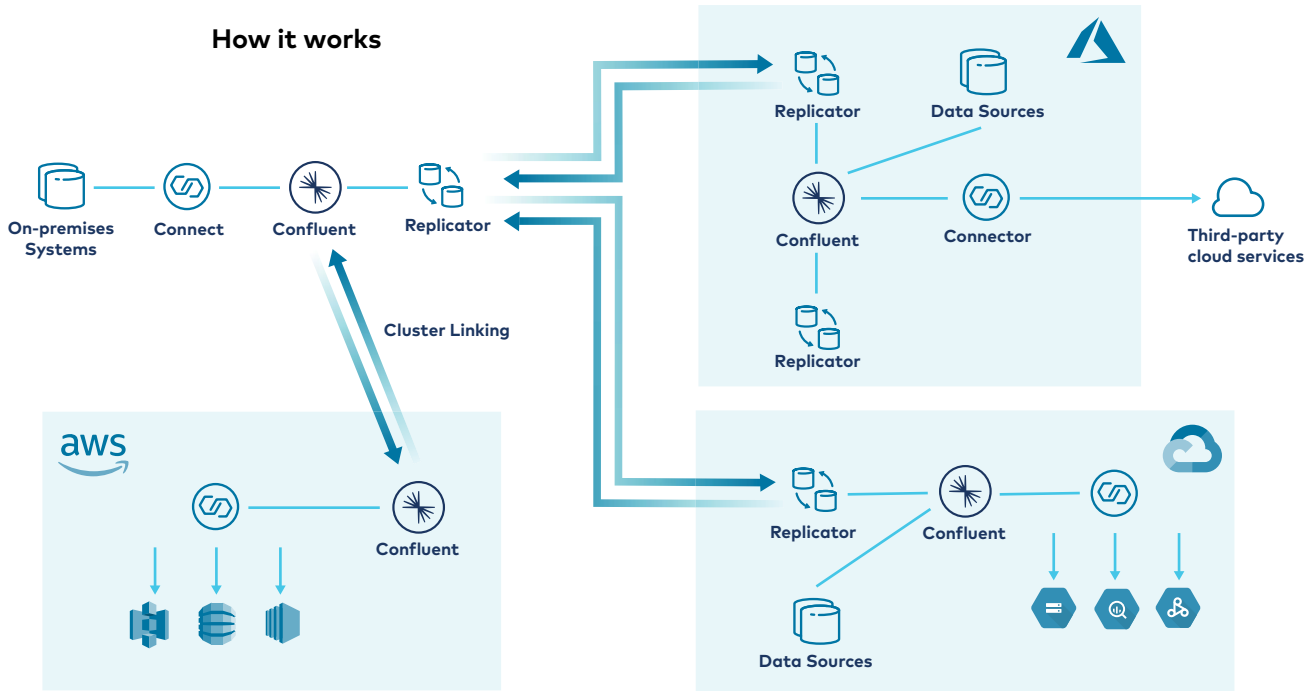
Oftentimes initial use cases are lower risk and simpler, like development projects, analytic workloads, and cloud backups. In trading we commonly see hybrid cloud use cases for things like risk calculation, where the simplicity and cost savings from cloud-native analytic services are often preferable to expensive on-prem HPC grids. Over time these use cases typically evolve into higher value production workloads where cloud services add more value, like training machine learning (ML) and artificial intelligence (AI) models to feed production trading systems or running production systems themselves in the cloud.

Whatever the mix, a hybrid architecture enables data to flow seamlessly to and from all the systems that need it, wherever those systems live. Many financial services organizations are building “data highways” using Confluent as the pavement. Data is continuously replicated from on-prem systems to the cloud, then fanned out to other applications, data services, and datastores using Confluent’s fully managed cloud service, [Confluent Cloud](#). Confluent is deployed on both sides of a hybrid architecture, seamlessly unifying on-prem and cloud environments and enabling data to flow to and from any systems that produce or consume data, globally.

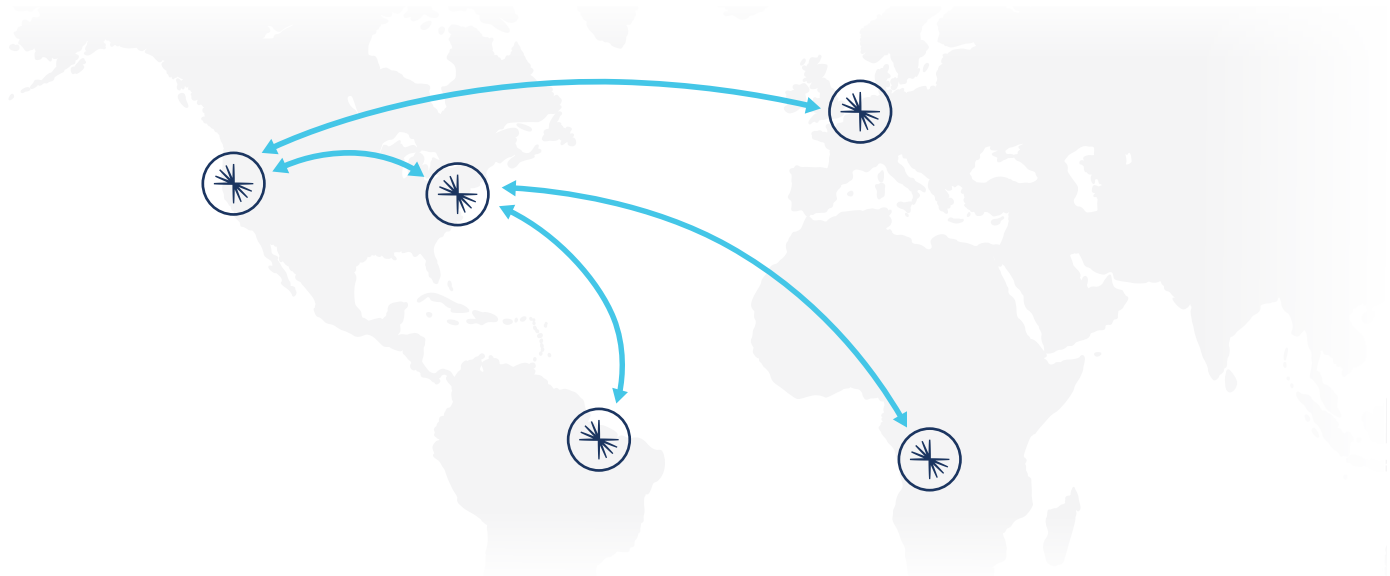
Additional benefits of a hybrid and multi-cloud architecture include:

- Ability to power legacy system migration from on-prem to cloud-native services in a phased approach, over time, with zero downtime or data loss
- Access to otherwise siloed data in on-prem systems for use by cloud applications; this empowers developers to choose the best cloud services for each project
- Enables multi-region and cloud disaster recovery
- Guarantees data consistency between systems and across global regions

- Connects edge data to centralized environments for IoT and edge use cases
- Manage geo-location of data to comply with regulations
- Cost savings from reduced operational overhead, reduction of legacy systems and licenses, and increased cloud provider discounts negotiated using multi-cloud optionality as leverage



To illustrate how a hybrid or multi-cloud architecture can add value for trading applications, imagine a scenario where all systems globally need access to the same data at the same time, like the price of a security, for example. Fixed-income traders in London may wish to access New York muni-bonds. A Confluent deployment in New York that contains this bond price data can be linked to a Confluent deployment in London, so that data is consistent across regions, enabling all traders globally to view consistent pricing with minimal latency and zero data loss. This capability unifies data silos within asset class technologies and across geographies to keep pace with the demands of the front office.





# Trading use cases: What your peers are doing

## Citi: Real-time dashboards for equities trading



Citi uses Confluent to join real-time tick data with reference data showing current positions, then display results in the UI for consumption by Citi's trading team. This is one of many real-time analytics applications for which Citi uses Confluent. See this [2020 Kafka Summit keynote](#) to hear Leon Stiel, Director at Citigroup, discuss how Citi is tackling its data challenges and using data in motion to drive efficiencies and improve customer experiences.

## Major market data provider: Real-time market data analysis

This customer built their real-time market data analytics platform on Confluent. They leveraged connectors to integrate their existing systems into Confluent, then processed millions of events per second using Kafka Streams, all with zero data loss and while meeting latency requirements. They use the Kafka Streams Processor API to build pipelines capable of handling millions of market movements per second with ultra-low latency. Kafka Streams is used to detect outliers, scoring (source confidence evaluation), arbitrage detection, and other processing of data in motion.

## Top 20 global bank: Order management system modernization

The client's proprietary trade order management system needed a technology refresh to provide greater throughput, resiliency, and fault tolerance. The new platform had to meet the following requirements:

- Process a minimum of 250 orders per second
- Scale linearly as data volumes increased
- Replicate data across datacenters in real time
- Recover and replay messages from a specific point in time (in case of unavailability of external systems)
- Guarantee message sequence for all events related to the same trade order
- No single points of failure

### Partner involvement

Risk Focus, a certified Confluent partner, worked closely with Confluent to develop an extensive fault-testing plan for the client's application while the Confluent team helped deploy and optimize the core infrastructure used to power the customer's OMS.

### Results

Risk Focus and Confluent successfully delivered the entire system on budget and according to the client's specifications. As a result, the client now has an application that can handle orders for its complete product line in a highly available, fault-tolerant system. For more information, read the [case study here](#).

## Top 5 global bank: Clearing automation and trade visibility

Strategic buy-side clients are demanding faster and more transparent access to information related to their trading business. This global bank uses Confluent in support of the following two areas:

### 1. NAV oversight:

- Buy-side clients demand a more real-time understanding of fund valuation and trade workflow
- NAV oversight provides real-time validation of funds coming from multiple systems and sources
- Clients can see where the trade is in the trading workflow, dependencies impacting the trade, and what's next

### 2. Asset services:

- Trade exceptions create operational risk and drive inefficiencies across counterparties
- Asset services provide a real-time exceptions dashboard to quickly identify open exceptions and associated risks
- Delivers an actionable prioritization of trades that need user attention before settlement
- Tie systems together to automate the clearing process and reduce the FTE overhead of manual processes, allowing operational staff to focus on exceptions

## Top 3 global bank: Global risk

BCBS 239 is a massive regulation that requires G-SIBs to aggregate and report on systemic risk across a variety of dimensions including, and not limited to, credit exposure, counterparty, liquidity, market, and operational risks. The breadth of this regulation requires integration across lines of business, various products, and the underlying IT systems that support them. This results in a requirement to integrate the right data in a timely and reliable manner in order to power reporting applications. Increasingly, G-SIBs are turning to Confluent and our data-in-motion platform to unify disparate silos of risk information locked away in databases, files, mainframes, and data warehouses. Confluent is deployed at one of the largest G-SIBs in the U.S. to collect risk information, in real time, from 40 different disparate risk and P&L systems and aggregate them into a singular data lake for BCBS 239 reporting. To accomplish this kind of data integration and movement without a centralized data-in-motion platform would have been prohibitively complex, expensive, and risky.

## Top 10 global bank: OATS and CAT processing

A large U.S. bank captures two billion messages per day from their equities trading systems using Confluent's JMS connectors for Solace and TIBCO EMS. This data flows through Confluent and into the customer's data lake for reporting. The OATS/CAT reports are then generated and sent to FINRA. This bank was able to retire their legacy data warehouse that was based on files and batch processing and, more importantly, set themselves up for intra-day reporting by continuously sourcing data directly from the upstream trading systems using Confluent. The scalability, durability, and resiliency of Confluent have set the bank up to add additional data sources to enable regulatory reporting for other asset classes, and new analytics use cases.

## Top 5 global bank: Drop Copy for real-time risk

Brokers often use a market's native Drop Copy services as part of their production trading operation to reconcile the trading activity reported by Drop Copy with the trading activity reported over production trading sessions in real time. A large broker uses Confluent to capture trade messages from their multiple OMS systems via both TIBCO and Solace and reconcile those with execution reports coming from NYSE, Nasdaq, and other equities markets to provide improved risk management. Confluent was chosen because we offer persistence, scalability, and performant stream processing capabilities using Kafka Streams to facilitate reconciliation. This helps the broker with real-time risk protection to avoid extreme scenarios where an issue occurs at the exchange or within the broker's production trading pipelines.

## Set your data in motion with Confluent

Data in motion is more powerful than data at rest. Trading systems and applications simply cannot afford to operate on stale data or be inhibited by inaccessible data that is locked away in disparate silos across global lines of business. Confluent enables trading businesses to access and utilize a combination of data and motion and data at rest together to power real-time applications that redefine modern trading software.

The trading systems that generate the most profit and most effectively mitigate risk must be fueled by a continuous flow of real-time data from all of the other systems, applications, and data sources they care about. This is nearly impossible without a centralized data-in-motion platform, and nobody has more experience in the trading ecosystem partnering with financial institutions to implement this kind of platform than Confluent. All 10 of the top 10 U.S. banks partner with Confluent to set their data in motion because Confluent is:

- **Cloud native** Confluent delivers true cloud functionality for data in motion. We offer a fully managed, cloud-native service that is massively scalable, elastic, secure, and globally interconnected, enabling agile development. This is a completely different experience than what would result from taking on-prem software and simply offering it on cloud virtual machines. With Confluent, developers and enterprises alike can focus on their applications and drive value without worrying about the operational overhead of managing data infrastructure.
- **Complete** We created a complete platform for data in motion by leveraging capabilities from open source Apache Kafka® with our significant proprietary capabilities. Our technology moves and processes data concurrently, with specific tools such as ksqlDB, a native data-in-motion database that allows users to build data-in-motion applications using just a few SQL statements, as well as over 100 connectors.
- **Everywhere** We have built a truly hybrid and multi-cloud offering. We can support customers in their cloud and multi-cloud environments, on premises, or a combination of both. From early on, we recognized that the journey to the cloud is not overnight or simple, and in order for our customers to effectively digitally transform, they require a fundamental platform for data in motion that can integrate seamlessly across their entire technology environment. This ability to let customers embrace the new without having to fully replace everything that is old is a critical point of differentiation and a critical element in the cloud adoption strategy of many of our customers.

The world is changing. Businesses are moving from a universe in which they *use* software to one in which they *are* software. Trading businesses who neglect this paradigm shift will be displaced by fintech startups with software development in their DNA, who are working tirelessly to disrupt incumbent players with better software products and services. This paradigm shift is already underway and it is accelerating quickly.

If you have any current or upcoming projects that might benefit from the capabilities listed above, we'd love to discuss them with you and see how we can help. [Contact us](#) now and somebody from our team will reach out right away.

## 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](http://www.confluent.io).