

FINANCIAL SERVICES

10 Ways Confluent Drives Transformation in Financial Firms

Overview	2
1. Consumer Banking	3
2. Payments	6
3. Fraud and Financial Crime	10
4. Capital Markets: Trading Operations	13
5. Risk Management	17
6. Insurance Transformation	21
7. Wealth Management	23
8. FinTech	25
9. The 21st Century CFO	28
10. Distributed Ledger Technology & Digital Currencies	31
Glossary	34

Overview

10 Ways Confluent Drives Transformation in Financial Firms

Right now, every aspect of the financial services industry is undergoing some kind of transformation. Consumers are demanding better services and better incorporation of technology. They are used to the easy access of social media, streaming services, and mobile devices, and expect that level of sophistication from the financial services firms that provide banking and investment services. Why shouldn't they?

It's not just consumers who are placing higher demands on the industry. All financial services firms face increasing scrutiny from regulators in every aspect of business, from how they manage their capital to how they describe their services and communicate with customers. Technology teams have very strict legal responsibilities defining how they manage data and how they keep it safe. They must provide a robust level of governance around the information they hold. Even the protocols that manage payments and trades are being upgraded to provide better security and improved standards of data privacy.

Regulations governing operational resilience demand that firms prioritize their most critical business services, give them the highest level of supervision, and put controls in place to keep information free from cyber crime and fraud. This is critical for running robust services, but can slow innovation. How can long-established businesses keep pace with young and nimble start-ups that don't have the same level of scrutiny or infrastructure to manage?

Financial firms have been encouraged by tech firms to move to the cloud for over a decade, but it isn't always as easy as it sounds, and the regulators are making sure that things are done in a comprehensive way. The cloud has a lot to offer: lower costs, dynamic scalability, reliable performance, comprehensive security, and easy integration. There are many incentives to persuade firms to jump through the regulatory hoops to get it done.

Data is the life-blood that helps firms bring all of their services together in the cloud. Firms need to be able to consolidate data from a variety of applications and services to be able to answer questions from the business. Data needs to be accurate, complete, and timely. People can't make decisions based on old or inaccurate data.

Software developers are building modern applications using micro services: small parcels of code that temporarily run in the cloud just long enough to do their job before disappearing. This new way of developing software is liberating for developers, but has a very high dependency on good data. Developers of microservices need to make sure that their microservice doesn't inadvertently take another service down if there's a problem.

Fortunately, Confluent solves a lot of these challenges by delivering a data platform that works in real time but is also decoupled from the applications that publish to or subscribe to data from it. This gives developers the confidence to write applications that can run reliably and the flexibility to add or remove services as they please, without risk to other developers. Customers benefit by using applications that always contain the most recent data, giving them confidence in their decision-making.

This whitepaper describes some of the financial businesses that rely on Confluent and some of the game-changing business outcomes that can be realized by using it.

1. Consumer Banking

Challenger banks and FinTech services pose the first credible threat to banking in a generation. Data supremacy is the key to success for both.

In order to survive the threats brought about by technology and innovation, traditional banking is being forced to re-think its core offerings.

Twenty years ago, internet banking was just beginning, and the first wave of online services were born. They were limited in functionality, tricky to use, and not very secure. Most consumers still visited branches to deposit cash and checks, and withdraw foreign currency. We took cash out of ATMs and paid for things using a card, cash, or check. Wire transfers were used for major purchases like houses or cars. If you owed money to your friends, you gave them cash. If you made a payment with a check, it would take three days to clear.

The State of Traditional Banking Today

Fast forward 20 years to our current ubiquitous internet, broadband, and mobile devices. Everybody is online almost all of the time, even in a plane or on the subway, and payments are instantaneous in most places.

Traditional banks today have changed significantly, but many branches have not evolved.

- A lot of branches have closed.
- Those that remain are being re-designed in favor of an automated experience with ATMs and IDMs (Integrated Deposit Machines). Automating the manual tasks frees up staff for more advice-related activities such as mortgage and loan consultations.
- ATMs still exist, but there are fewer of them in production, and they are notoriously unreliable.
- ATMs are also a significant target for criminals; not just a target for being vandalized, but a convenient host for skimming devices designed to read customer card details. Many fraud attacks involving debit cards originate from ATMs.
- Telephone banking exists, but is costly to operate and less popular with consumers.
- Most services are available online, but innovation has been painfully slow. A global financial crisis didn't help.
- Most banks have mobile apps now, but they don't yet have all the services customers expect.

Up-and-Coming Rivals: FinTech and Challenger Banks

Specialized FinTech firms provide a wide range of financial services from personal payments to foreign exchange, lending, and insurance. The business models deployed by FinTech firms have disrupted the entire financial services sector, and forced the industry to re-think how it can remain competitive. Traditional financial firms are being forced to transform their entire business models, not just their apps.

In the last 10 years, a variety of challenger banks have emerged, offering services that resemble a conventional bank but without the branch network, ATMs, or expensive central office. Challenger banks compete by offering a better online and app experience than consumers have been offered by their existing bank. Services have been built from the ground up, are cloud-native, and have been designed with user experience in mind. They are free of legacy infrastructure like mainframes to slow them down. Large technology companies like Apple, Amazon, and Google are also entering the consumer financial services space. Many are acquiring banking licenses and have plans to become insurance carriers, too.

There are many reasons for banks to be fearful, because these new services are often very good. But traditional banks still have one thing on their side: lending. Banks make the bulk of their profits through lending money, either through mortgages, which are relatively low risk, or through personal and commercial lending, which is higher risk but a rich source of profits. FinTechs also lend money. They generally don't have the historical data they need to be great at credit risk, but they soon will, and are turning to new techniques like machine-learning models to try and accelerate their progress. So the race is on. (FinTech is a huge topic, and described in another section of this whitepaper.)

How Confluent Helps Banks Remain Competitive and FinTechs Profitable

As discussed above, customer experience is critical in the delivery of all consumer-facing financial services. This includes:

- Ergonomically elegant design of all aspects of tech-based service, from the look and feel down to the navigation and search. Consumer apps are the benchmarks that need to be equaled.
- Services need to perform quickly, with minimal waiting times.
- Outages are simply unacceptable.
- Payments need to be transacted in real time.
- Information needs to be presented in a clear and understandable way.
- Significant work needs to be done to make offers relevant, timely, and financially suitable for the customer.
- Firms targeting customers with investment products or loans need to meet stringent regulatory requirements.

All of this requires data supremacy. Firms need to have all of the relevant data on hand, updated in real time, and available to be used on demand. Old data is useless. Here are three examples of how Confluent prepares companies for success: setting up the basics; customer state vector technology; and high-quality, constantly updated risk data.

The Basics

Delivering the level of service today's customers demand requires a mastery of the basics, including:

- Robust IT infrastructure
- Comprehensive security controls
- A fraud prevention system in place
- Instantaneous payments
- Customer confidence in responsible management of personal data and identity information
- Only when all of this is taken care of can firms start to deliver an advanced customer experience. Confluent starts by taking care of the basics.

Delivering Hyper-Personalized Customer Experiences

Instacart, an e-commerce platform and Confluent customer, maintains a real-time customer state vector system. They keep a record of each customer that is constantly updated and informs them of the best possible journey for each customer interaction. Applied to financial services, this concept can be used to match customers to products and services and serve up the best possible content almost all of the time. Confluent delivers this game-changing technique based on real-time data stored and maintained for every customer. Every time a new data point is available it is instantly added to the vector, ensuring that customers are always shown the best possible content.

High-quality Constantly Updated Risk Data

Credit risk is the most critical of the risk disciplines that must be mastered in consumer financial services. Credit risk protects a financial firm from future credit losses and defaults, and enables a firm to optimize the efficient frontier of rates with losses. If credit risk is incorrectly measured, rates are either too high or too low for a specific customer or basket of customers. It's a constant balancing act, because every day loans mature and new customers must be found, and credit teams need to build a balanced portfolio of credits with different credit qualities and potential earnings.

What has any of this got to do with data? Well, in order to be good at credit risk, the traditional approach is to run models that are trained with historical data. This has been the approach of banks for years. FinTech newcomers don't have the luxury of long data histories, but do have modern data skills. Many have built machine-learning models that can challenge the effectiveness of the old methods without needing so much data.

Managing credit risk requires high-quality recently updated data and the ability to react quickly to consumer requirements. Confluent is the real-time data platform that can provide that foundation.

Confluent Outcomes in Consumer Banking

The state of traditional banking is under threat. FinTech and challenger banks represent daunting competition. Confluent can help by providing:

- Better customer experience
- Faster, more agile credit risk
- Better fraud prevention
- Better security
- Real-time payments
- Robust controls over personal data
- Reduced reliance on mainframes, and subsequently lower costs
- Accurate and timely regulatory reporting
- Secured and highly performing Open banking APIs (APIs are application programming interfaces, which enable developers to create software applications.)



You have real-time data: stock price, things that are ticking. You have data that's pretty static: terms and conditions, things like that. And then you have data that's updating periodically: things like position updates. If you can use a tool like Apache Kafka® to pull all that data together and combine it in ways that you display to end users—whether they be traders, salespeople, managers—and you provide them analytics across that data, it's extremely powerful.

— LEON STIEL, DIRECTOR, CITIGROUP.

View the complete story [here](#).

2. Payments

Hyper-growth and innovation against a background of cyber threats.

A second arena for Confluent is within the dynamic payments industry, which is going through a revolution alongside traditional banking. Payments technology has been embedded into almost all consumer services, from e-commerce to gaming, transportation, television, and mobile apps. Consumers have a multitude of choices about how and when they pay and a range of devices that make it easier to do so. Payments apps make it easy to split a bill for dinner or make a group booking for a concert, and the adoption of these services has been faster than the industry expected.

At the same time, the payments industry is bringing in new standards for messaging, security, and fraud prevention as well as tightening up rules and procedures. There is a lot for all the participants in the process to work on, and significant gains to be realized from making the systems more robust.

The payments industry has exploded due to:

- Technology improvements leading to new online payment platforms and services.
- The growth of real-time payment services in most regions.
- The rise of embedded payments in e-commerce, gaming, and leisure apps.
- The availability of APIs to embed payments in software applications.
- The demise of branch banking and ATMs.
- The reluctance to use cash in some markets, accelerated by the Covid pandemic.
- The ease of using mobile devices to pay for goods and services in stores, on public transport, and just about everywhere else.

Now that we have these technical advances and increased demands on the global payments infrastructure, the organizations that run the different components of the global payments ecosystem all need to scale up their operations. This hasn't been an easy journey for everyone. Meanwhile, banks and consumer-facing digital services need to meet the rising demands of consumers, tougher and arguably better regulation, and an increasingly challenging security environment. Customer experience needs to be best-in-class to compete.

New Demands on Payments Services

To stay competitive, payment services must integrate significant upgrades in:

- Customer experience
- Performance and resilience
- Real-time services, e.g., Faster Payments / SEPA, and online payments
- Cyber security
- Fraud and financial crime
- Increased standards and regulation, e.g. ISO 20022
- Compliance factors like PCI
- Open Banking and APIs
- Embedded payments services

Data in motion. In recent years, there has been a major shift in the way firms think about how data is processed and managed. Firms are moving away from relying on traditional systems of storing data-at-rest to thinking about how their businesses can react to events in the real world as and when they occur. This gives them the agility to deal with changing business requirements, effectively managing data in motion. In the payments industry, this is an incomplete journey; many countries and systems have real-time payments capabilities while others have been slower to adapt and are still working on the solution.

How Confluent Transforms the Payments Industry

Confluent has become the central nervous system of many of the major participants in the global payments ecosystem. This includes networks, gateways, online payments platforms, and banks.

- It has been particularly popular with FinTech firms launching new cloud-native platforms, as they depend on real-time access to data to support their microservices-based applications.
- Confluent connects all the payments networks that banks have relationships with to the bank's core infrastructure, using adaptors and APIs. This enables aggregations to be built to support multiple views of different payment rails.
- Confluent delivers an aggregate view of all of the bank's payments activities, regardless of how or where they originate.
- Confluent is a real-time data platform, meaning that it tracks and retains all the key payments data as the payments are received, processed, and settled. Payments can be kept in Confluent briefly or indefinitely; it's up to the individual company.
- Whilst payments are normally streamed into the bank's mainframe, having the data available in Confluent in real time means that it is available for leveraging into a broad range of relevant applications.
- This platform approach also takes the load—and critically, the cost—away from the mainframe while making data that would otherwise be hard to obtain available to business-critical services like fraud detection and security operations.

Working with Mainframes

Mainframes are still used extensively by banks and payment networks, largely for historical reasons, but they persist due to their reliability and the cost of changing to an alternative technology. Using Confluent to 'offload' critical data from the mainframe allows firms to make it accessible to any system that can take advantage of it within seconds. This includes next-generation cloud applications, such as microservices. Importantly, this does not require a disruptive 'rip and replace' process; it allows the mainframe to remain the system of record while Confluent runs as parallel architecture that can take on new workloads over time, reducing cost and risk.

Open Banking and PSD2

Open banking is a framework that allows third-party firms to build applications and services around many financial institutions using APIs. It has placed additional demands on the way banks manage payments.

The European Union passed PSD2, the Payments Services Directive, to encourage innovation and protect customers in a single regulatory effort. It made it necessary to be able to initiate a payment using a set of open APIs and created yet another interface banks must secure and protect from fraudulent activity. Service Level Agreements (SLAs) must be met and customer activity tracked.

Cyber Security

Fraud has become a growing problem in payments and is harder to detect than before because of the increasing array of channels, platforms, and networks that present themselves as a wider attack-surface to cyber criminals. Firms use Confluent to conduct real-time fraud detection that in many cases is able to block a transaction before it can be completed.

The goal is for firms to be able to identify fraud in real time, and notify the customer about a suspected fraudulent transaction instantly, avoiding the need to try to recover funds or process the claim with a customer. This is a major cost saving, as firms often spend \$3 for every \$1 of fraud to process the claim and reimburse the customer. (See the separate section on fraud and financial crime.)

In recent years, cyber security has become front of mind for the payments industry, and high-profile incidents (such as the attack on the Bank of Bangladesh specifically targeting the Swift network) have highlighted the need to monitor all aspects of the payments process for suspicious activity. Meanwhile, IT operations teams have had to work very closely with their counterparts in security to harden their networks and software applications and test for vulnerabilities. When firms need to use machine-learning models for advanced security investigations, Confluent is used to feed the models real-time data to train and deploy the models rapidly (sometimes on the stream) and make the intelligence part of an operational business process. That allows firms to act on intelligence in real time, often while the threat is still active.

Money Laundering

Money laundering regulations dictate that banks conduct rigorous checks on payments that are being made, checking watchlists provided by governments and regulatory authorities. They also have to run equivalent sanctions checks, blocking any payments that don't meet the correct criteria. Confluent enables these watchlists to be checked in real time, meaning that compliance doesn't add any latency to the transaction.

Payments Networks

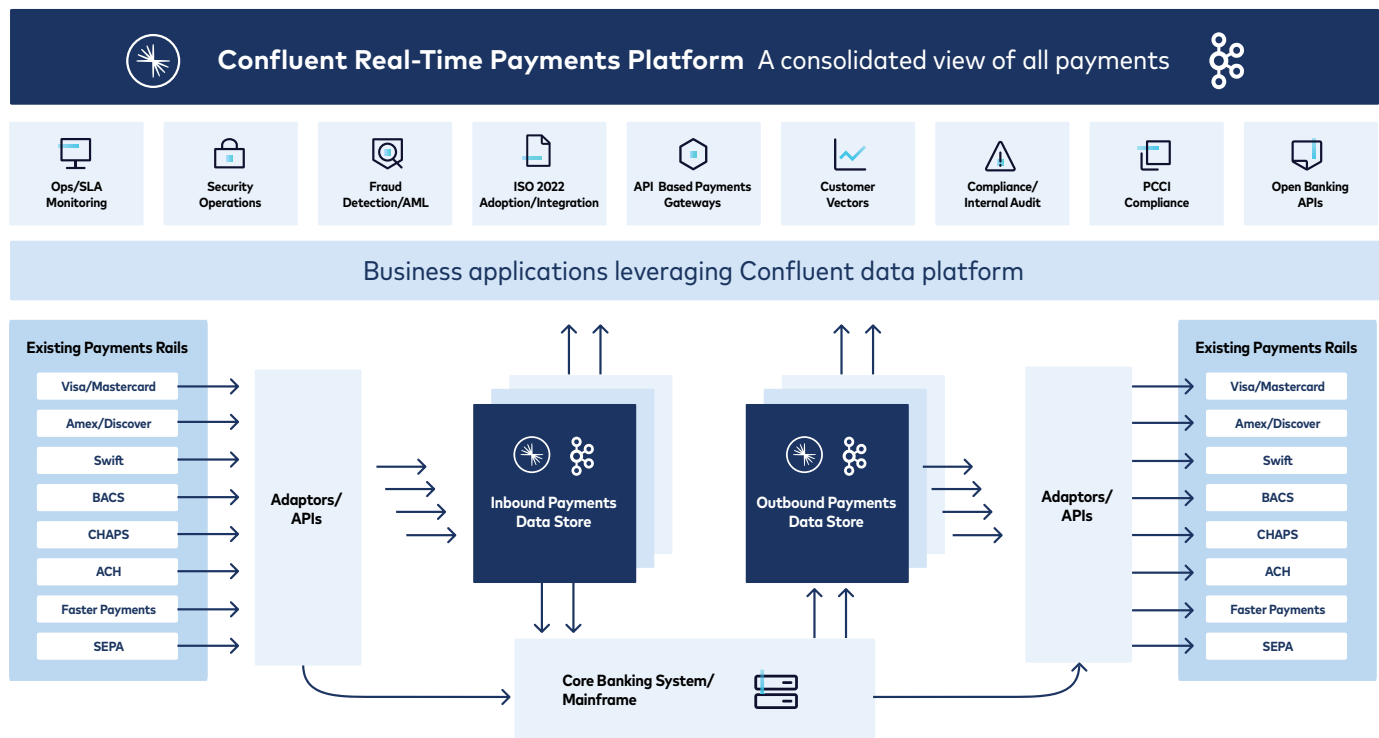
Many of the same challenges faced by the banks are present for payments networks. There is a focus on performance, scalability, and resilience. Networks are dealing with different stages of the payments process and have a different view of the data; they aggregate data using different criteria. Networks have to deal with a multitude of payment standards and messages formatted in different ways with different characteristics. They also have to manage relationships with banks, gateways, and internet services in different regions, managing payments services with different SLAs and operating with different legal requirements and regional data jurisdictions.

Confluent is used to enable data integrations and transformations in real time on streaming data, allowing networks to cope with the multitude of different formats on both inbound and outbound payments within strict SLAs without significantly slowing down the process. This set of capabilities makes it possible to transition to the requirements of ISO 20022, an open global standard for financial information. Confluent allows this new operating model to be adopted and cope with the additional data contained in the payload of the payments message and the operational changes required to support the regulation. This includes the single-event structure facilitated by the new data model, making the process fully event-driven.

Online Payments Services

Online payments services share similar challenges and have to design their systems to work securely while embedded in third party applications like e-commerce websites, streaming services, or gaming platforms. APIs need to be secure, identities need to be monitored, and multi-factor identification solutions need to be incorporated. Volumes can be very high and subject to seasonal fluctuations, which is where the cloud-native elastic scaling of Confluent provides stability.

The diagram below shows how Confluent is used as a persistence layer for payments with robust APIs that can drive multiple business applications from a single view of the data. For data to persist means the data survives after the process that created it is over.



Confluent Outcomes in Payments

One of the most powerful aspects of deploying Confluent in a Payments environment is the ability to take a platform approach: bringing the core payments data into Confluent and using it to drive an array of valuable applications that all utilize a common set of data.

Payments firms using Confluent all benefit from these common outcomes:

- Improved customer experience through managing customer state vectors.
- Higher performance, better resilience, and elastic scalability.
- Security improvements.
- Real-time fraud detection and prevention.
- Money laundering prevention and sanctions compliance.
- ISO 20022 and PCI compliance.
- Easy access to information through a stable platform
- The ability to launch new services faster.
- Reduced reliance on mainframes and legacy systems.
- Lower operations costs
- Transaction tracing—using Open Telemetry standards.
- Real-time alerting of key events related to a payment.
- Faster time to market
- Legacy modernization: the ability to build the new system alongside the existing data infrastructure and with the same data.



Capital One delivers risk insights in real time with stream processing.

Capital One supports interactions with real-time streaming transactional data using Apache Kafka. Kafka helps deliver information to internal operation teams and bank tellers to assist with assessing risk and protect customers in a myriad of ways.

Inside the bank, Kafka allows Capital One to build a real-time system that takes advantage of modern data and cloud technologies without exposing customers to unnecessary data breaches, or violating privacy regulations. These examples demonstrate how a streaming platform enables Capital One to act on their visions faster and in a more scalable way through the Kafka solution, helping establish Capital One as an innovator in the banking space.

View the full talk [here](#).

3. Fraud and Financial Crime

Moving from detection to prevention by acting on better information in real time.

Firms have been able to run some aspects of fraud detection in real time for many years, but speed has always come at the expense of quality.

Quants (quantitative analysts) have spent years trying to optimize the relationship between fraud detection and false positives.

Accurate detections have been too low and false positives have been too high. It's not surprising, really, since most transactions come with very little circumstantial evidence to guide a model or an investigator in their decision-making.

Most transactions find their way into a mainframe and detection systems rely on static statistical models and rules for detection. Some have evolved to machine learning or AI, but are only as good as the data with which they are fed, which is somewhat limited. Since firms first set out to solve this problem, there has been exponential growth in channels of transaction and in the volume and complexity of the attacks. As such, the number of indicators that can influence a decision has grown, and the approach required for each indicator can be subtly different. Indicators can vary from the location of a person to the size or currency of a transaction. Some fraud attempts are very simple to spot, whilst others require a more sophisticated approach.

Since the 2020 pandemic, fraud teams have had to adapt, as many fraud-detection models needed to be adjusted for the new normal of working remotely and shifting consumer behavior. Firms had to create a different baseline for 'normal' behavior. During this period, fraud incidents spiked to new highs. Once many people returned to the office, things changed again, although not to the normal we used to model for.

The boundaries between cyber security and fraud/financial crime have blurred. Cyber attacks on all aspects of financial services are often the first stage in fraud taking place. Frauds no longer just happen in seconds; some of them can last for years. Consider common attacks like phishing, or account takeovers. Are they cyber attacks, fraud, or both?

Cyber attacks and fraud are too closely linked to be considered independently. The problem is that many firms have structured their teams independently, along with the systems and processes that support them. Things need to change in order for firms to get on top of the threat. Not only are these domains closely linked, but the bad actors are changing their tactics daily. Teams need to be agile in their approach and have a portfolio of techniques to call upon.

As the result of a major crime that took place in Canada in 2016, firms are being advised by both regulators and their own auditors to consider building a fusion center, a place where cyber security and fraud teams merge their operations and form a new team dedicated to fighting all aspects of cyber security and financial crime.

It makes a lot of sense. From the data perspective, fraud teams have typically managed transactional data originating from the general ledger or databases. The data is good, but mostly related to the transaction or to personal data about the customers. Conversely, security teams have access to vast quantities of unstructured data that originate from the logs produced by websites, software applications, hardware, authentication systems, etc. Individually considered, these are useful data sources. Combined, they are like gold—for crime fighters, anyway!

Enabling Fusion Centers

Collaboration is the most important thing required to make a fusion center successful. Firms need executive sponsorship and a willingness to break down the barriers between teams and their data silos. Political fiefdoms must be broken down. Teams need to embrace a new, collaborative way of working, often with data and processes that will be new to them. Data is the crucial component.

Firms use Confluent as the foundation of the Fusion Center by bringing together transactional data from the ledger, mainframe, or databases, including watchlists and external data, and joining it with all of the unstructured data sources found in security operations (which can run into the thousands). This would have been a daunting challenge using conventional technology, but with Confluent's platform approach and the ability to bring one data source at a time on board, it becomes possible to achieve big wins early in a project.

The most powerful form of detection requires rich contextual data around a transaction. Simply knowing that a person in London wired \$50 to a person in New York isn't enough to determine if that transaction is legitimate or fraudulent, even if you know who they are and

have a history of their transactions. Bringing in contextual information from website logs can tell an investigator the IP address of the user and potentially their location and device type. Authentication records can tell you how often they log in, where from, and if they have had unsuccessful attempts to log in before. Records of user behavior show when a user has changed their password or added a new recipient for funds. All of these events combined might reveal that the person from London is normally based in Germany, made the transaction from a new laptop using a new type of browser, and recently changed their password before adding a new recipient in New York. This is more than enough information to block the account and ask for authentication prior to allowing the transaction to take place. This kind of intervention reduces fraudulent transactions and reduces the cost of remediation after the event.

Contrast this with the context that requires firms to spend \$3 for every \$1 of fraud committed. Clearing up after the event requires expensive resources. Moving from detection to prevention improves customer experience and saves everyone money and time.

Confluent customers say they can alert their customers of suspicious activity on their cards “before the customer’s card is back in their wallet”—i.e., just a few seconds from tapping their card in a store to a text message warning. Not only does this reduce fraud, but it improves the customer experience by reassuring them that their card issuer is looking out for them.

Confluent’s use within fraud detection is widespread and is used by many different types of businesses engaged in many activities. Some examples include:

- Account takeover (applicable in banking, cards, payments, and eCommerce)
- Transaction Fraud (applicable to all businesses engaged in digital transactions)
- Wire Transfer Fraud (applicable to firms wiring money across any network)
- ATM Fraud (including skimming attacks, jackpotting, and proximity attacks)
- Phishing, smishing, spear-phishing
- Money laundering and sanctions violations
- Gift-card scams
- Man-in-the-middle attacks
- Invoice fraud
- Insurance premiums fraud
- Insurance claims fraud
- Market abuse
- Rogue-trader incidents

Confluent Outcomes in Fraud Prevention

- Moving from detection to prevention by acting on information in real time.
- Bringing together fraud/financial crime and cyber security operations in fusion centers.
- A single set of data to fight fraud, cyber crime, money laundering, and all other types of digital financial crimes.
- Improved customer experience by preventing fraud.
- Substantial cost savings through avoided fraud losses and the associated downstream costs.
- Higher retention rates from customers who feel they are better protected.
- Reduced downtime of systems taken down through fraudulent activity.



Royal Bank of Canada (RBC)

Streaming Data Empowers Royal Bank of Canada to be a Data-Driven Organization

Challenge

Improve the customer experience through data while staying compliant with industry regulations

Solution

Implemented Confluent Platform to stream data in real-time, migrate to an event-based architecture and embark on a microservices transformation

Results

- Higher quality, faster and scalable customer services
- Increased efficiency in application building
- Lowered anomaly detection time from weeks to seconds
- Implemented data reuse across teams for relevant business insights
- Greatly reduced cost

"There isn't another product that competes with Confluent Platform. Streaming data as events enables completely new ways for solving problems at scale."

— MIKE KROLNICK, HEAD OF ENGINEERING, ENTERPRISE CLOUD, RBC

View the full story [here](#).

4. Capital Markets: Trading Operations

Continuous innovation in a market full of challenges and risk.

Capital markets have seen significant reforms in the last decade with robust risk management, closer supervision, new asset classes, and technical transformation. The regulatory environment has become more complex, some regulations have come and gone, and reporting has become more rigorous and frequent.

Capabilities that seemed impossible only 10 years ago can now be delivered using recent technologies, although businesses have often had to adapt and modernize processes to keep up with the technology.

The decision-making agility that firms now require dictates that they have a very broad range of information sources readily available to them, from relatively static reference data and internal sources to streaming feeds from market data providers and exchanges.

Firms have to cope with a highly complex environment:

- A combination of complex systems running in local data centers.
- Their own hardware being operated in co-located facilities provided by trading venues for low-latency operations.
- Some services run in the cloud, with the need to bridge to on-premise applications.
- Some services are delivered by third parties.
- Network performance is mission-critical.
- The ability to aggregate data from all remote operations is essential to manage positions and portfolios and control capital and risk effectively.
- Many firms have operations in over 30 trading centers, with co-located operations in as many trading venues.
- New requirements to trade on cryptocurrency exchanges.

It gets even more convoluted. In modern capital markets, traditional assets are traded alongside cryptocurrencies, which are evolving on a daily basis. Conventional participants are trading alongside retail speculators, low-latency firms, hedge funds, and other agile participants, not all of whom are subject to equal regulation. The rush to compete in cryptocurrency has led to the development of numerous competing crypto products, stablecoins, and Central Bank Digital Currencies (CBDCs), the rules for which are still being determined. All of this has a huge impact on capital markets firms, who need to move fast to keep on top of developments and capitalize on them.

This hybrid approach to software and infrastructure requires a modern approach to data management. Immediate answers don't come from data at rest. Firms have had to learn to work with streaming data pipelines to get instant answers, and treat databases as systems of record, using them for historical data analysis.

How Confluent Enables Trading Transformation

Confluent has become the central nervous system of trading operations data for many of the world's largest brokerage firms, hedge funds, asset managers, and exchanges.

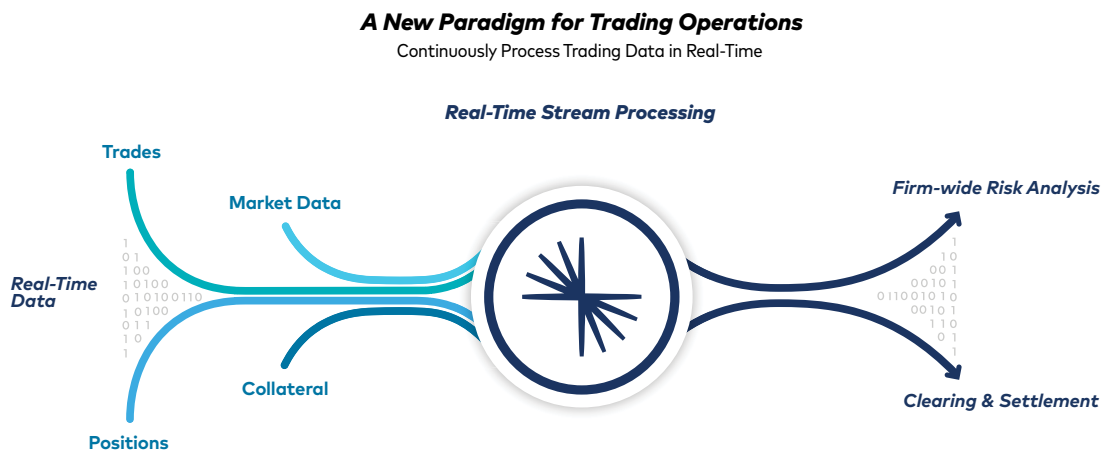
Firms have a multitude of inbound and outbound data sources that must be captured and transformed, ready to populate the systems that the business uses to make decisions.

Something as simple as an individual trader maintaining an accurate positions table requires every trade to be loaded and the positions calculated and updated in real time. While this sounds simple enough, consider a firm with thousands of trading staff, thousands of instruments, thousands of portfolios, millions of positions, and millions of trades, that trades in 30+ cities in multiple time zones using multiple currencies.

Confluent breaks down this complexity by enabling aggregations to be created rapidly at a level of detail appropriate to the task, at the appropriate level of the hierarchy. Streaming data keeps these aggregations continuously up to date, and the people who need to be informed have the correct information readily available.

Trading teams take this aggregated information and enrich it with other sources of information to conduct a multitude of tasks. Tasks vary, but include:

- Portfolio valuation
- Calculating exposures relative to limits
- Comparing critical ratios like hit and fill ratios across venues
- Monitoring roundtrip times
- Optimizing collateral
- Managing margins
- Trying to stay within the guardrails set by risk teams in a risk-appetite framework



Other trading teams build algorithms to deploy in their low-latency environment. These algorithms require rigorous testing and take full advantage of the replay capabilities within Confluent that allow them to replay historical scenarios to see how the models would perform under a range of conditions. Monitoring these models is critical to making sure they perform as expected. When the models deteriorate, they must be stopped and re-calibrated. Confluent is critical to this entire process.

Algorithm developers use this same technique to measure the performance of any model either in or out of production, meaning that they are constantly aware of which models would perform the best at any point in time under a certain set of conditions. This capability allows them to rapidly switch models in production to suit prevailing market conditions.

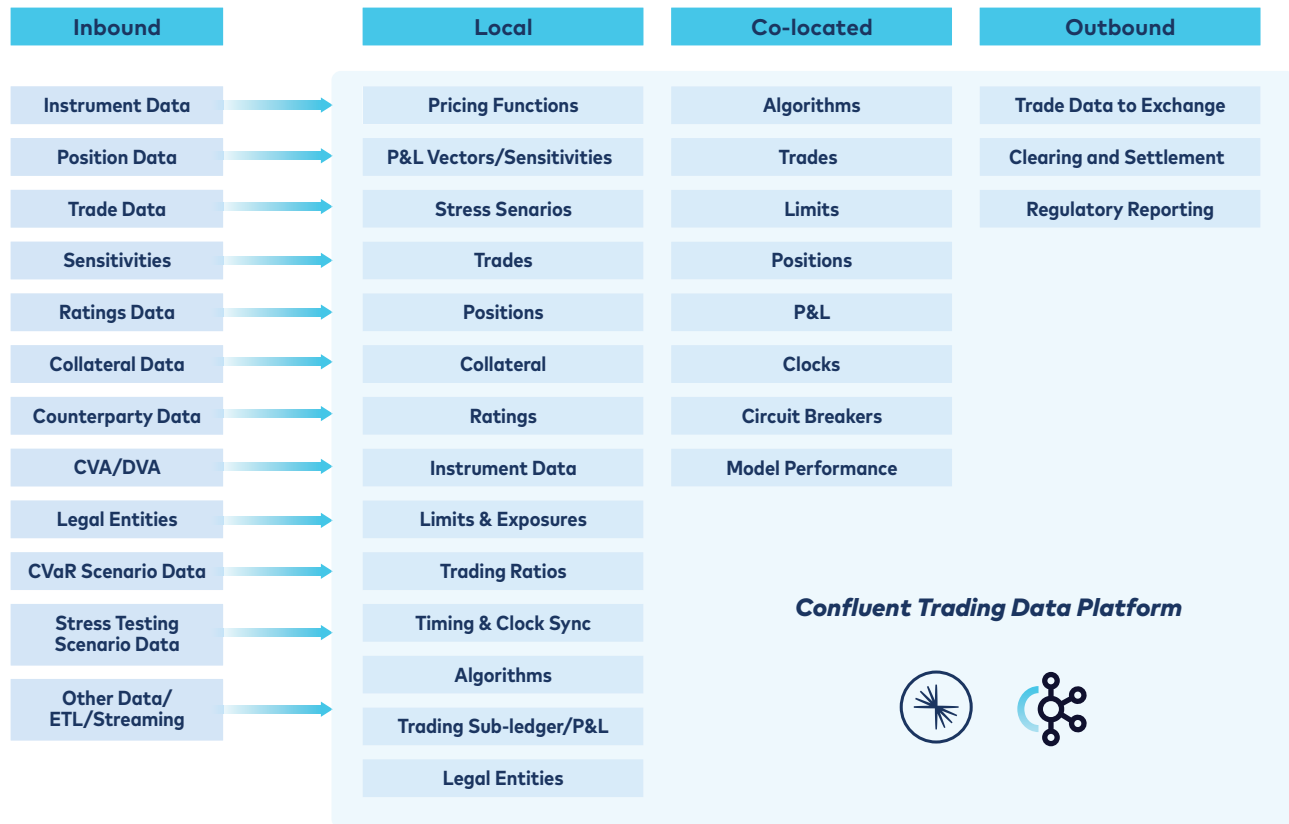
Canceled and amended trades are a constant feature of trading environments. Trades often don't complete for legitimate reasons, like a lack of market liquidity or a sudden change in prices. Sometimes trades partially fill but then orders expire when the market closes, leaving traders to reconcile the balance with finance. This is a normal part of trading operations, but it is also a great source of information for investigators looking at the trading book for sources of bad practice or rogue trading activity. In the past, bad actors could take advantage of this report only being produced once a day; so long as they covered up their activity by the end of the day, people wouldn't find out. Not any more: real-time canceled and amended trades reporting closes that loophole.

Finance teams require their own aggregated view of the data to manage the trading sub-ledger and P&L at all levels of the trading hierarchy, along with capital ratios, collateral, liquidity, daily reports, and other critical measures.

Risk teams need the same information to calculate market risk, counterparty credit risk, CVA, DVA, liquidity risk, model risk, CVaR, and to calculate and maintain sensitivities (P&L Vectors). Risk is a critical part of trading, but is a large topic and is discussed in the next section.

There are multiple technical use cases for deploying Confluent in post-trade processing, settlement, and clearing. This varies by market and by the different participants involved in the process, but typically it is used to create a persistence layer for trades from which a number of other processes can be conducted. Firms have used Confluent as a staging environment between their different trading systems and settlement and clearing. It ensures a high level of consistency and quality in the data.

The diagram below shows some of the data sources critical to trading.



Confluent Outcomes in Trading Operations

- The ability to stream any data from any system and build real-time views of operations.
- Aggregating data in any combination and at any level of the hierarchy.
- Creating position tables by joining streams
- Maintaining a persistent record of positions at a portfolio level or other aggregations.
- Measuring key trading statistics like hit and fill ratios.
- Enabling algorithm developers to test their algorithms on replayed historical scenarios, and monitor the algorithms against live market data.
- Building real-time "canceled and amended trade" reports.

- Delivering a real-time view of the trading sub-ledger to enable finance teams to measure P&L and other critical finance measures.
- Delivering a firm-wide view of risk across all asset classes and all risk measures in all regions in real time.
- Measuring critical technical performance indicators like roundtrip times and venue seasonality to inform smart order-routing algorithms.



Capital Markets: Trading Operations, Euronext

Pan-European Stock Exchange Relies on Confluent to Power Event Driven Trading Platform

As the first pan-European exchange—spanning Belgium, France, Ireland, the Netherlands, Portugal and the UK—Euronext operates regulated securities and derivatives markets in Amsterdam, Brussels, Lisbon and Paris, as well as a regulated securities market in Ireland and the UK. Euronext recently developed a new event-driven trading platform, Optiq®, that provides a tenfold increase in capacity and an average performance latency of as low as 15 micro-seconds for order roundtrip as well as for market data.

Underpinning the Optiq platform is a persistence layer that the Euronext development organization built using Confluent Platform. Confluent Platform provides a reliable, scalable streaming infrastructure for Optiq that supports millisecond latencies with no messages lost. "When we started, Kafka was a new technology to us, and one that we had decided to use for a very critical application in our system," says Philippe Planchon, Architect and Innovative Trading Solutions Director at Euronext. "With Confluent we felt supported in our decision and we knew we had the right level of expertise to get prepared and to help if we encountered any issues. That was a key element in our success."

Challenge

Develop a new trading platform for markets across multiple European countries that supports high-volume, high-speed trading and provides clients with access to real-time data.

Solution

Use Confluent Platform to implement a reliable, scalable persistence layer for market orders that supports millisecond latencies and billions of messages per day.

Results

- Reliable 24/5 operations achieved and maintained
- Stringent performance requirements exceeded
- Dedicated, expert support received

View the full case study [here](#)

5. Risk Management

Delivering a real-time, firm-wide view of risk; meeting the regulators armed with robust information; teaming with the CFO to plan effectively

Risk management teams have found themselves firmly in the spotlight since the 2008 financial crisis. Since that time, techniques evolved, new asset classes emerged, and the data used to drive risk calculations became more readily available and higher in volume and complexity.

During the same period, risk teams have grown. All firms now have a Chief Risk Officer, and many CROs now have their own dedicated technical staff (CRO/IT). Regulations have also evolved, becoming tougher in many cases, and in some circumstances subsequently being relaxed. Billions of dollars have been spent on risk projects, driving consulting firms and IT vendors to specialize in the subject.

Risk is core to a firm's strategy and is fundamental to achieving profitability. Firms that master risk and bake it into their operations will be better-run, more stable businesses, with better prospects than those who do not.

How Confluent Helps Firms Manage Risk

Firstly, firms need a multitude of information sources in order to build a risk-management platform. Many of these sources are common to trading, as demonstrated in the previous section, but they only deal with the data required for trading in capital markets. Risk management is also undertaken by asset managers, hedge funds, retail and commercial banks, insurance and reinsurance firms, and anyone else who puts capital at risk. They all have their own data challenges, regulations, risk measures, and reporting requirements. Despite this, many tools and techniques are used in common. We will focus here on managing trading risks in capital markets, but rest assured that Confluent is equally applicable to managing risk in all financial markets.

What are the data requirements?

Firms need access to all the sources mentioned in the last section for trading, plus many more: notably the data sets that support stress testing, which are covered later in this section. Much of this data has to be gathered from multiple global operations; one recent project involved streaming data from 37 different trading locations.

This data has to be ingested in high volumes and aggregated at every level of the business hierarchy, making each aggregated set available to be used by risk engines that calculate a variety of risk measures and calculations, and also for risk managers who interpret the results.

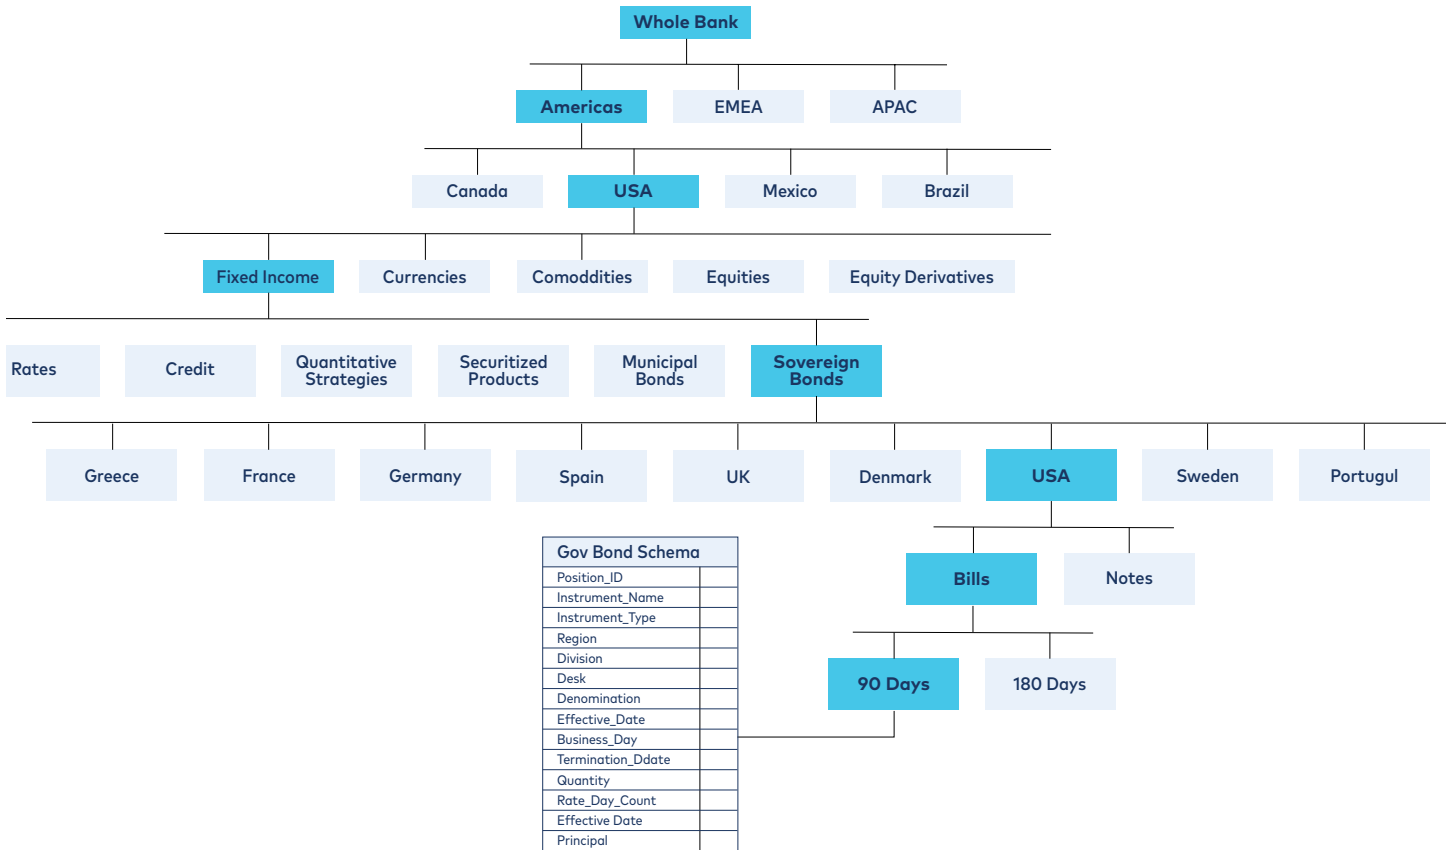
The risk types calculated include market risk, counterparty credit risk, CVA, DVA, liquidity risk, model risk, CVaR, and the calculation and maintenance of sensitivities (P&L Vectors). These measures represent significant complexity, and demand accuracy in data throughout the risk process. Streaming data keeps these aggregations continuously up to date. To allow risk teams to continuously measure risk at the appropriate level of the hierarchy, each level has to be calculated separately, as many of the risk measures used are non-additive. The stream of data enables a team to build a firm-wide view of risk. At the time of the global crisis in 2008, this was done daily, and leaders often had to wait a full day to understand their exposures, so having the information on hand in near real time represents huge progress.

When errors are made and incorrect data is loaded into risk models, there needs to be a robust reconciliation process to correct the incorrect data and reload it into the various systems where it is required. Teams must take great care to make sure that records are not accidentally duplicated in the process of reloading, and serialization techniques must be applied. This is a good match to the capabilities of Confluent, as it can make use of the AVRO serialization framework to make sure that events are stored in exactly the order in which they occurred. This case is especially relevant in low-latency trading environments where thousands of trades can occur in a single second. It is imperative to know what order they took place in. Accurate serialization informs the logic that ensures events are not loaded multiple times erroneously.

This requirement becomes even more relevant when firms have to deal with intraday risk reporting, which requires separate streams of data to lock down their risk situation at a given point in time, enabling them to give an accurate report of specific risks at a specified point in time as the business continues to operate. This is made more complex by the fact that not all of the information arrives at the same time due to different latencies of different inbound data sources. Latency could vary from fractions of a second to many minutes.

Once an accurate set of data has been established, firms are able to take the outputs of all of these risk systems and build a risk war-room for the CRO, who can see all of the relevant risk information in context and take measures to manage risk at a firm-wide level.

The diagram below shows an example of a typical business hierarchy in a trading environment.



Stress Testing

Risk teams also have to prepare regulatory stress tests for different purposes. These range from the portfolio level right up to whole-firm stress tests, and the team must be ready to carry out regulatory fire drills, which happen on a regular basis. Confluent enables these stress tests to be conducted rapidly, taking some mathematical problems that would normally be calculated in batch jobs and makes them continuous. Techniques like streaming Monte Carlo simulations become possible, enabling some aspects of stress testing to be conducted in real time.

Passing or failing a stress test determines whether or not a firm can trade in the markets the following day. It is crucial that firms conduct their own tests regularly so that they can be sure to pass the tests when they are set by regulators.

In the preparation of stress testing scenarios, risk managers have to design a range of scenarios from normal to extreme trading conditions. The objective is to identify the conditions and combinations of conditions that will break the firm. Being armed with this information allows them to do a better job of setting asset class strategy, limits, and exposures, and in some circumstances changing the risk profile of the firm. The scenarios have to be applied to the current positions in the trading book and be run thousands of times with multiple combinations of various stress conditions. Each run must be analyzed after the event. This has historically taken a long time to complete; often several days. Confluent can reduce this process significantly, with some parts of the process yielding immediate results. Multiple days of work can be reduced to minutes using this approach. (It is worth noting that this also applies to consumer banks, who have to stress test all of their lending books, and to insurance carriers, who have to stress their entire book of business under the Solvency II capital requirements.)

The diagram below shows an example of the sort of data and processes that are required in a risk department, and would provide the required information to risk teams to enable them to conduct stress tests.



Confluent Outcomes in Risk Management

- A flexible data platform for risk management, allowing risk teams to pivot rapidly and conduct new analysis with zero notice.
- Streaming aggregations at all levels of the business hierarchy with the ability to create new hierarchies on the fly.
- Positions and portfolios are continuously updated in real time from streamed trade data and market prices.
- The measurement of multiple types of risk from a single set of data.
- A firm-wide view of risk with sufficient information for CROs to make rapid decisions on a firm's trading strategy.
- Rapid stress testing of scenarios and the ability to run a whole-firm stress test on demand.
- Creation of scheduled and ad-hoc regulatory reports to comply with current and future regulations.
- Stream lineage: a full audit trail of every stream, mapping the relationships between data points and hence improving data governance and visibility.
- Identifying and fixing errors in data processing as they occur as opposed to interrupting lengthy batch processing schedules which then require a restart.



Nationwide Building Society (UK)

Challenge

Embark on a digital transformation to respond to competitive market as well as new banking regulations, increase digital activation, respond to increasing payment volumes, and meet member expectations for customer experience.

Solution

Investing in technology such as Confluent and digital capabilities to be primed for the next generation of digital innovation. Confluent is the foundation for an "event hub" and part of a resiliency and agility program.

Results

- Confluent is enabling agility and autonomy in digital development teams
- The first use case with Speed Layer should remove c7bn requests / year from the HPNS
- Kafka will help Nationwide maintain its service availability despite unprecedented demand
- Kafka and streaming being adopted across multiple lines of business, both top down strategic use cases, but also bottom-up engineering led demand
- Speed Layer, Streaming and Kafka will help Nationwide head-off the threat from agile challenger banks

6. Insurance

What does a modern insurer need to be competitive in a changing risk landscape with rapid technological innovations?

Insurance firms have a complex juggling act to maintain. Balls in the air include pricing risk effectively, winning new business through an increasing array of channels, maintaining a balanced portfolio of premiums, and marketing effectively in a highly competitive environment, all in a world with increasing regulatory requirements.

Advances in technology have forced firms to rethink the way they manage information. Tools like telematics have shown great innovation and flexibility, allowing more young drivers to get insured. Driverless vehicles, while not yet legal on the road, are forcing the issue of how to deal with different types of risk and the need to monitor customers in a way that has not been considered before. Issues of ethics and law meet those of monitoring real-time streams of information and understanding the risk of a potentially large portfolio of new premiums for an untested line of insurance.

Meanwhile, threats to the business persist. Cyber security and financial crime present constantly shifting challenges, particularly with a more agile and remote workforce. The 2020 pandemic highlighted the agility of the industry including a change in the distribution of premiums underwritten and the requirement to re-think the way premiums for travel and health are priced and sold.

Aggregators continue to win business on price, and competition remains fierce across all the different channels to market. New InsureTech players have entered the market intending to disrupt the industry in the same way that has been experienced in other areas of financial services. Brokers, aggregators, and carriers have to work harder to keep their services competitive. Those who have the ability to use information creatively can use it to full advantage when seeking to win new business and retain the business that they fought so hard to win the first time around. Now that the loyalty premium gap has been closed in many countries, firms are going to have to deploy new tactics to maintain margins.

How Confluent Transforms the Insurance Industry

Competition is stronger than ever in the insurance market, and the added threat from nimble new entrants is present. Firms can do a lot to compete by making sure that the customer experience is stronger than ever before. This starts with data. Making sure that there is a seamless and consistent experience across channels is a given. Firms can compete by making sure they have excellent data on their customer base and prospects. This data can be used to maintain a customer state vector for each prospect and customer and make best efforts to target the right offer to the right customer at the right time. Many firms are pursuing the multi-policy route across many lines of business; this has the potential to be very profitable, so long as data is managed well and information on existing premiums in-force is used to inform premiums pricing for new and multi-line cover.

Firms continue to want to renew their customers' policies; however, they should consider all of their options at renewal in regions where consistent pricing is now enforced. Pricing remains a useful way to discourage customers whose risk profiles have deteriorated and are no longer desirable. Multi-line cover, however, can be very profitable for the correct customer profile. These decisions can only be made well with accurate, up-to-date data.

Managing risk is still at the heart of insurance for actuarial teams, for those managing the Solvency II capital model (or equivalents), and for the risk management teams in the asset management side of the business.

Risk management is covered in a previous section. However, Confluent is highly applicable in areas of the Solvency II capital model, stress testing scenarios, and in market, counterparty, and credit risk within asset management. All of these areas require a robust data framework of timely and accurate data. Having this platform in place leads to faster report generation and the ability to flex to ad-hoc reports when necessary.

Underwriters have to be more agile than ever in order to ensure that premiums pricing remains accurate in rapidly changing market conditions. Teams need to be able to forecast accurately, and to do so need current data on premiums in force, lapsed premiums, fraud, and information from external sources, including economic data. Models should be monitored closely and calibrated regularly. Claims and

the processes for managing them need to be streamlined. The time from FNOL (first notification of loss) to settlement often determines the size of the final settlement. Firms make sure everything is done to settle a claim promptly. Claims inflation is common, as is fraud, so statistical techniques and sometimes machine learning need to be used to detect and minimize outliers. Brokers or agents have a part to play here, and the information exchange between brokers, agents, other third parties, and carriers needs to be robust.

All of these areas require new technologies to be embraced. Telematics, IOT Sensors, and mobile applications are all helping insurers evolve from the expensive legacy systems that many still support. Most big firms still have mainframes and need a way to be able to work with them whilst supporting modern technology systems. Confluent allows firms to connect to the mainframe, and provides a persistent source of information from the mainframe that can drive modern application development frameworks like microservices. This enables companies to take full advantage of emerging technologies without incurring additional mainframe costs or physically replacing the mainframe.

Confluent Outcomes in Insurance

Insurers deploying Confluent as their data platform can expect to be more agile in many aspects of their business, including:

- **Underwriting:** access to a broad library of data to support pricing decisions.
- **Risk:** a data platform to facilitate daily execution of pricing functions, ad-hoc risk modeling, capital model calculations, and the execution of stress-testing scenarios. Risk will be significantly more agile.
- **Asset Management:** the ability to manage portfolios more effectively backed up by near real-time data, faster risk information, easier reporting, and more timely valuations.
- **Claims:** shorten the claims process by bringing all of the information together and conducting fraud and other checks faster.
- **Cyber Security Operations:** acting on new threats faster and enabling security operations teams to feed models with fresh and accurate data.
- **Technology:** easily implement new technologies like telematics and IOT sensors while working with old technology like mainframes.
- **Fraud Prevention:** improving detection in the premiums application process and in claims.
- **Customer Experience:** using customer state vectors to maintain a real time picture of a customer and their situation.

7. Wealth Management

Delivering superior advice from superior information.

The wealth management industry is also undergoing a transformation, although perhaps not as quickly as some predicted. Automated advice has its place and is a growing market for people with smaller portfolios, but there are still significant upsides to paying for bespoke advice from humans. Despite improvements in technology, this is unlikely to change any time soon. Strategies on how to deal with tax, inflation, and succession planning are too personal and complex to automate effectively.

The wealth management industry has to be able to support both forms of advice and leverage the information superiority of automated advice to deliver a superior experience to private clients. Technology platforms can help determine which of their clients and prospects will need help with specific issues at times determined by the markets and the global economic situation. Armed with better data, wealth managers will deliver a better customer experience by being relevant and timely in their communications and in the products they recommend.

How Confluent Helps Wealth Management Firms Deliver Better Service

The challenge is that wealth management firms have to deal with a lot of complexity in their operations. Data is the key. Firms operate in multiple regions, supporting clients with complex portfolios that span asset classes, geographies, tax laws, currencies, and competing priorities. Affluent individuals need complex tax strategies and often borrow in multiple regions to spread their risk and exposures.

Wealth managers need to understand which products work well in different markets and determine which clients to target and when. Strategies developed for some clients won't work for others. At the high end of the market, advisors need to offer a highly personalized service; while for the mass affluent, there are many more potential clients and so advisors need to decide where and who to invest their time in.

This is where data can really help.

The leaders of wealth management firms need to take a holistic look at their base of clients and prospects and decide where to focus their attention. This drives decisions on hiring, marketing, and product design. Confluent helps by bringing data together from CRM, finance, and HR systems, plus products and market data to build a platform that answers all of those questions.

The same platform can be used to drive client-facing applications that allow clients to view their portfolios and communicate with their advisors. Confluent helps to update portfolios automatically by continuously keeping position tables updated every time there is a trade or a corporate action. This means advisors have an accurate view of customer portfolios at all times. (See the section on trading for more details.)

From a data perspective, aggregations are critical. Product managers need to understand what works where. Advisors and their leadership in distribution need to understand who is being effective and why. Which strategies are proving popular? What leads to customers churning? Often the questions are quite simple, but the data required to give a comprehensive answer is harder to come by. Aggregating real-time data from finance and distribution enables profitability attribution to be calculated.

As in other branches of financial services, customer information is critical. Confluent enables firms to maintain a real-time customer state vector that gives an advisor enough information to communicate the right information at the most appropriate time for each customer—also referred to as “next best action.”

This is marketing gold. If advisors can prioritize their workloads based on data, they will advise clients more effectively, improving relationships and thus their NPS (net promoter score).

Advisor retention is also critically important. It takes years for advisors to build good client relationships, which can be rapidly lost when an advisor leaves. Wealth firms can retain their best advisors by monitoring critical indicators linked to advisor behavior. Simple measures like the number of meetings conducted, products recommended, products on-boarded, new funds invested, and NPS can all help identify the top performers who deserve extra retainment efforts.

Confluent Outcomes in Wealth Management

- Organizes data from CRM, Products, Risk, Distribution, and Finance to reduce complexity and make analytical decision-making easier and faster.
- Updates portfolios and assesses risk continuously.
- Continuously updates customer state vectors to drive customer information delivery and enhanced customer experience by delivering the next best action.
- Utilizes continuously updated machine learning to make recommendations faster and more precisely.
- Creates aggregations and calculations on the fly, enabling comprehensive profitability attribution.
- Allows advisors to act quickly on market conditions based on an agile decision-making platform.

Affin Hwang Asset Management Launches Digitization Initiative

With more than RM 80 billion (\$19 billion USD) under management and close to 200,000 clients, Affin Hwang Asset Management Berhad is among the leading asset management firms in Malaysia. The firm's roots stretch back to 2001, and over the past two decades its IT infrastructure expanded organically as it grew to meet the needs of institutions, pension funds, government-linked companies, and other clients. As a result, the back-end landscape included multiple siloed legacy systems and batch processes that, although fully functional, were beginning to strain under the load placed on them and affect the customer experience.

To address this challenge and modernize its operations, Affin Hwang Asset Management launched a digitization initiative that is being powered in part by real-time data streaming with Confluent. Confluent is enabling Affin Hwang Asset Management to integrate siloed systems, reduce batch processing, and explore new product offerings.

Challenge

Modernize and digitize core operations at one of the largest asset management firms in Malaysia

Solution

Use Confluent to replace slow batch-oriented processes with data streams that provide near-real-time updates to key systems

Results

- Time to market cut by months
- Day-long data access delays eliminated
- New business opportunities opened

"At Affin Hwang Asset Management, we are great believers in not wasting time, so once we started our digitization initiative, we wanted to start making progress right away without long, drawn-out setup times for new technology," says Allen Woo, Chief Innovation Officer.

"The team at Confluent has been exceptional in understanding the resources that we have and working closely with us to support our vision of adopting data streaming quickly. Already, transaction data updates that used to take up to a day are now available right away, and we're excited about the many other doors that data streaming will open for us."

Although Affin Hwang Asset Management is in the early stages of its event streaming journey, the company is on a fast track to maximizing value by putting more of its data in motion. "We'll be streaming onboarding and payment data next with Confluent, and we're looking at incorporating streaming data with AWS Step Functions and Lambdas," says Woo. "In addition, we have more training sessions planned for our teams with Confluent engineers, which will help us continue our adoption in an agile manner."

8. FinTech

The not-so-quiet revolution.

Over the past 10 to 15 years, a revolution has been taking place within the financial services industry. A wave of startup companies have arisen with the simple intention of performing one or more of the financial services that consumers need with modern tools, zero legacy, and a fresh way of thinking about problems. This uninhibited approach has led to a wave of innovation and disruption. Not everyone has succeeded, but many have, and a lot of them have done so in a way that is both profitable and delivers a significantly better customer experience than those that came before.

No part of the financial ecosystem is safe from disruption. Things started with payments platforms like PayPal initially facilitating peer-to-peer payments and then rapidly broadening into embedded payments for consumer websites and growing into a global payments platform. Others followed, and there is now a competitive market in payments services, supporting peer-to-peer payments, embedded services, cross-border payments, micro-payments, and merchant services. Now anybody with a coffee cart and a mobile device can set themselves up to accept payments, making businesses more agile and cost-effective. We are already seeing a second wave of disruption in which a new wave of disruptors have hit the market and the original disruptors are seeing their business models challenged.

Many conventional firms have disrupted themselves and are competing effectively. Those who didn't are losing business rapidly.

At the same time as the payments industry was being transformed, other firms emerged to offer a range of services, including (at the time of writing) FX services, lending, regulatory technology (RegTech), insurance quote aggregation (InsureTech), wealth management (WealthTech), pensions, eCommerce funding solutions, account aggregation apps, and challenger banks.

During this transformational period many things happened, not least the introduction of the Payment Services Directive 2 (PSD2). This set of regulations, which started life in the EU, has given rise to what is known as Open Banking. The regulations stipulate that banks must publish (at least) two sets of APIs, allowing FinTech firms and other competing banks access to view critical data about an account holder, and the ability to initiate payments into or out of an account.

This capability is a goldmine for FinTech services. The ability to access bank accounts has stimulated the banking industry to be competitive with the offerings that their smaller, more agile competitors are building. It means that a whole new class of FinTech has emerged, offering account aggregation services where a person can use a single app to see their aggregate funds across all of their bank accounts and other financial services. This is great for consumers, as it makes it easy to move cash to the place with the best rates or offers.

Conversely, it is a honeypot for cyber-criminals and fraudsters who have a large new attack surface to go after, with the potential reward of a direct entrance to a bank's customers. Banks have had to invest significantly in making these new interfaces secure.

Challenger banks have been around for a while now. Early entrants like Monzo and Starling Bank in the UK have inspired a flood of new contenders in many parts of the world. Supporting technologies like faster payments, contactless payments, and 5G have enabled these new competitors to deliver superior service. Many of these firms struggled in the early days with customers using them as a secondary bank account, taking advantage of the service quality while continuing to rely on the relative safety of their conventional bank. This, however, is beginning to change as consumers turn to challenger banks for loans, and many younger people are choosing these accounts as their first and primary bank account. As challenger banks bolster their balance sheets and pioneer new credit risk techniques, they should see a steady rise in prominence, and conventional banks will have to disrupt themselves to remain relevant.

Something that is common to just about all of these FinTech firms is the ability to deliver a level of customer experience equivalent to that experienced by consumers using social media or entertainment apps. This is something that sets them apart from legacy competitors and builds brand loyalty. Many social media platforms and entertainment apps rely on Confluent to measure and deliver their customer journeys. This is an important topic that we cover in more detail later.

One area of FinTech that is experiencing rapid growth is Distributed Ledger Technology (DLT).

Distributed Ledger Technology has created a whole new class of disruption with the ability to prove the authenticity of events, secure transactions, facilitate smart contracts NFTs, and enable new currencies, from crypto-currencies (Bitcoin, Ethereum, etc) to Central Bank Digital Currencies, some being pegged to a sovereign currency. This topic has its own section later.

How Confluent Transforms the FinTech Industry

Most FinTech firms interact with their customers exclusively through apps and websites/APIs. By default they are technology businesses and are intolerant of any loss of service or performance. Latency is as bad as downtime. Customer experience is key.

Most FinTechs build their services on state-of-the-art-technology platforms and are cloud-native by design. These technologies include Microservices, Kubernetes, and Kafka. They also make use of many SaaS platforms. It makes sense that firms are developing their applications using these technologies, as it means that they can scale elastically to meet peaks in demand and can be extremely resilient. They need to be built using robust techniques and methodologies that keep bad actors out of their networks and minimize exposure to security vulnerabilities and fraud. Many of the services need to connect to each other.

FinTechs widely deploy Confluent as the data backbone of their products. Microservices-based architectures are very demanding of the data they take advantage of and have many dependencies. The loosely-coupled nature of Confluent allows individual microservices to subscribe to data topics in Kafka and not be held back when other services are unavailable. Confluent also serves as a great way of joining home-grown applications with SaaS services.

The massive scalability of Confluent means that in periods of high demand, there are no issues with availability or performance.

Confluent Outcomes in FinTech

For FinTech firms, operational resilience is key and 100% service availability is taken for granted. Much of this is dependent on good design. Using Confluent as the central nervous system for their data has many benefits. They include:

- Elastic scalability, resilience, and fault tolerance.
- Flexibility that enables new features to be added quickly and other systems to be connected with ease.
- Loosely-coupled architecture, meaning that applications can cope well when other systems are experiencing latency.
- Excellence in customer experience, using Confluent to maintain customer-state vectors and using this knowledge to deliver a unique experience for every customer.
- The ability to build robust security into services.
- Prevention of security breaches and fraud/financial crime through vigorous monitoring of applications, data, and infrastructure.



NORD/LB Supports Multi-Year, Bank-Wide IT Transformation with Confluent

With total assets of more than 164 billion euros, NORD/LB Norddeutsche Landesbank is one of Germany's leading commercial banks. The bank is in the midst of an enterprise-wide transformation program, NORD/LB 2024, in which it will invest about 500 million euros in new IT systems and digitalisation initiatives.

As part of the NORD/LB 2024 program, the bank is developing a new core banking platform based on event streaming with Confluent. In addition to enabling more advanced, real-time analytics across the bank, this platform will support improvements in fraud detection, customer retention, and trade/investment analysis, as well as new products that will differentiate NORD/LB from its competitors in the market.

Challenge

Help drive a bank-wide transformation that includes a 500 million euro investment in new IT systems and digitalisation initiatives

Solution

Use Confluent's platform for data in motion to develop a new core banking platform, enabling improved realtime analytics, fraud detection, and customer retention

Results

- Competitive differentiator established
- Event streaming adoption streamlined
- Streaming infrastructure costs reduced

"Confluent is enabling us to address our need for a scalable, highly available messaging infrastructure that allows us to decouple our producers from consumers, setting the stage for us to be more flexible, agile, and responsive to change."

"Confluent's technology continues to provide the opportunity to develop competitively differentiated products for the future and is the technology we'll take our next steps with."

— SVEN WILBERT, NORD/LB

9. The 21st Century CFO

Modern businesses serve customer demands in real time. The modern CFO needs to be perfectly informed with a stream of real-time insights that drive decisions.

Not long ago, CFOs and the teams supporting them were masters of Excel and had to work with a plethora of tools to support the core processes of a finance department.

There were two major problems. Firstly, teams had to work really hard just to keep up with the business, manually wrangling data into the correct shape and spending more time building reports than working on the findings they contained. Secondly, most of the information they delivered would be on a daily or weekly basis, analyzed in a morning meeting over coffee.

For most people, that's just how it was, but the world has moved on and things have to be different.

In a modern real-time business the CFO needs an operations center—a mission-control style war-room where the finance teams track all the financial operations of a business in real time. Real-time businesses can change quickly, and sometimes the CFO needs to apply the brakes or step on the gas. Either way, that can only happen if their business telemetry is in real time. There are numerous teams and processes to consider within the remit of a finance team, and most of them benefit from a steady stream of reliable, timely data.

There are multiple ways in which CFOs are relied on by their peers to bring real-time insights into the daily running of a business. Examples include:

Capital Management: Hedging of foreign exchange and commodities.

Sales. It is critical to be able to take a real-time view of a sales pipeline, being able to examine the different stages of a sales pipeline, how key metrics are performing against a baseline, and how inbound leads are converting from marketing leads into sales leads, and then into fully qualified opportunities. Is the firm tracking against its expectations?

In financial firms, this takes on an additional level of complexity with credit pipeline forecasting. It looks the same as a conventional sales pipeline, but with the additional layer of credit quality. Forecasting credit quality is an essential component of the sales process. It is critical to a firm achieving a balanced portfolio of credit qualities. If credit quality is tracking down, why? What is the root cause of the change?

Also critical to sales teams is the requirement to pay sales compensation, and by default, the requirement to fund those commissions and forecast them accurately. Accurate forecasting is critical to avoiding surprises down the line.

Key business metrics need to be tracked, including a wide range of KPIs related to sales, retention, costs, and customer satisfaction.

Real-time information on business performance makes it possible to run promotions optimized in real-time and attuned to the specific needs of individual customers. This concept applies to all kinds of businesses, but for financial firms it means being able to deliver uniquely personal experiences to customers and deliver products that are perfectly suited to their needs.

How Confluent Helps CFOs Transform Their Organizations

In a financial firm things are even harder than in other industries because the core business is producing high-velocity financial data, whether that comes from payments, banking transactions, or the trading floor. CFOs have to work in unison with Risk and Compliance, and decision-making is collaborative. Regulations like iFRS9 (which tasks firms with forecasting credit losses, reporting the valuation of all assets accurately, and accounting for the changes in their valuation on the firm's P&L) force the different teams to work together and account for the risk-taking decisions of the business.

iFRS9 is a great example of a regulatory requirement that can be simplified with a data-streaming platform. Forecasting in this scenario requires the generation of multiple credit-loss curves for a broad range of scenarios. Having that data readily available with the ability to serve multiple scenarios shortens the process and allows more scenarios to be considered.

It doesn't end there. There are so many other regulations CFOs have to build their businesses around— some purely financial and others that require deep collaboration with their risk counterparts. Another notable example is BCBS 248, Intraday Liquidity Management. Like IFRS9, this was a consequence of the global financial crisis of 2008. Firms are required to manage and report on liquidity during the trading day, not just in overnight reports. BCBS 248 forces firms to change their operating model for risk, and have the ability to run a report on their liquidity positions in an on-demand fashion. To make it more challenging, firms have to be able to run ad-hoc, intra-day Liquidity Risk stress tests. This is non-trivial, and requires close collaboration between teams. (Stress testing is described in more detail in section 3.) To produce an ad-hoc stress test, the basic data on all the positions in the trading book need to be continuously updated, along with the sensitivities and all of the relevant instrument data and market prices.

Stress tests are only useful if the baseline data is correct. This can be significantly streamlined by being continuously prepared. One of the strengths of Confluent is being able to replay data at any point in time, and this capability can be leveraged to run a sequence of stress tests rapidly. Without a streaming data platform, this would be prohibitively time-consuming. Stress testing can be made less stressful.

Confluent transforms the way finance teams work by building a data platform that brings together data from all of the relevant systems and processes to one common platform. With Confluent, you are able to stream massive volumes of disparate financial data in a loosely coupled way. What this means is that a business doesn't have to build hard-coded integrations with all of their systems. Systems and data can change, and that's ok. Confluent gives teams the ability to subscribe to data streams and make that information available to a range of stakeholders in a way that is appropriate for their needs. Because the data is so easily available, it can deliver a wide variety of business applications. These include:

- Real-time, firmwide dashboards with accelerated telemetry for key functions.
- Data frameworks for regulatory calculations and reporting.
- P&L monitoring and attribution.
- Real-time promotions and offers.
- Collaborations between finance and other teams to solve a range of complex business challenges.



Bank Rakyat Indonesia

Bank Rakyat Indonesia Increases Financial Inclusion and Extends Microfinance Opportunities Using an Event-Driven Architecture Enabled by Confluent Platform

With 75 million customers and a market cap of more than \$38B, Bank Rakyat Indonesia (BRI) is the largest bank in Indonesia and the largest microfinance institution in the world. BRI is in the midst of a digital transformation that is simultaneously strengthening the bank's position among Southeast Asia's most successful financial institutions and enabling it to increase financial inclusion across Indonesia, which has one of the highest unbanked populations in the world.

The transformation at BRI is enabled by an event-driven architecture based on Confluent Platform and Apache Kafka that is powering real-time analytics for credit scoring, fraud detection, and merchant assessment services. Among the innovations being driven by event streaming are a system that detects anomalous customer transactions in real time, an ISO 27001 certified open API that connects BRI with a digital ecosystem of partners, an early warning system that identifies customers who are at risk of payment default, and a microlending app that is helping BRI reach untapped new market segments.

Challenge

Drive a digital transformation at the largest bank in Indonesia to improve the bank's market position and increase financial inclusion across the country

Requirements

Move from synchronous to asynchronous microservices development on an enterprise-ready platform. Enable stream processing for real-time data processing in flight.

Solution

Use Confluent Platform and Apache Kafka to deploy an event-driven microservices architecture that powers big data analytics for real-time credit scoring, fraud detection, and merchant assessment services.

Results

- Fraud detection performed in real time
- Loan disbursement times cut from two weeks to two minutes
- ISO-certified open API created
- Loan defaults predicted proactively; NPL at 0%

"Confluent Platform and Apache Kafka, by enabling us to build and deploy real-time event-driven systems for credit scoring, have helped BRI become the most profitable bank in Indonesia.

"We love open source, but at the same time we're not a startup. We're a large financial institution that works with world-class organizations, and we need services that make it easier for us to sleep at night. Confluent Platform is very reliable; it's never down. It has become our backbone."

– KASPAR SITUMORANG, EXECUTIVE VICE PRESIDENT AT BANK RAKYAT INDONESIA

10. Distributed Ledger Technology & Digital Currencies

A tale of two immutable logs

The Distributed Ledger Technology (DLT) space is little more than a decade old, but has already experienced multiple laps of the hype cycle. Formed officially with the release of Satoshi Nakamoto's Bitcoin white paper, the last 10 years have seen a trillion-dollar industry appear around cryptocurrencies, blockchains, closed consensus distributed ledgers, and smart contracts. As is the way with disruptive technologies, probably a few too many acronyms and buzzwords followed the actual innovation, so there's a glossary included at the bottom to help.

A distributed ledger is a consensus of replicated, shared, and synchronized digital data geographically spread across multiple sites, countries, or institutions. Unlike with a centralized database, there is no central administrator.

The development of DLT has the capacity to facilitate monumental change in how financial processes are handled. Most of the supporting systems were built around central authorities and humans policing financial procedures. Although moves have certainly been made to automate parts of these processes, their bottlenecks persist. DLT moves the obligation of regulation to cryptographic rules and smart contracts. When this smart contract becomes transactional law, it allows intermediaries like clearing houses to be removed from the transaction and streamline these processes to minutes rather than hours or days.

The distributed nature of DLT provides an additional market opportunity. Moving away from a traditional reliance on multiple ledgers not only simplifies reconciliation, but also exposes potential new abilities to bring finance to the next generation of the internet. This shift in mindset centers around decentralizing and opening knowledge resources; data can be shared and reused across applications, enterprise, and community boundaries without a gatekeeper. This means it is possible to not only offer new products on top of the new data source, but also interact with the protocol more efficiently. The former is clear in NFTs (Non-Fungible Tokens), where a unique key becomes the distributed virtualisation of a given asset like art or music, while the latter allows new approaches to how credit risk is calculated or cross-border payments are handled. This open network directly encourages application ecosystems that are continuous and asynchronous, a trend that persists more widely across the finance space.

Regardless of their astronomical growth, businesses in this space are not without their challenges. Despite the underlying ethos of decentralization, the average consumer or institution requires some degree of centralization during their experience to ease onboarding, accelerate adoption, reduce risk, and provide security and insurance. For example,

From this need for centralization within decentralization, an economy has been born. It offers the containing services that handle everything from KYT (Know Your Transaction) risk analysis to cryptocurrency exchanges and many more across a variety of market verticals. This provides a point of centralisation on the edge of a decentralized network. These services become a gateway for those without the capital, interest, or risk tolerance to become a node or wallet in the network itself.

These services and companies expose the digital asset economy and Decentralized Finance (DeFi) to traditional finance and other industries, which means they begin to adopt the challenges of those industries too. Some wished for these innovations to directly replace the world of TradFi, or traditional finance,, but the reality is that TradFi and DeFi have instead become measured alternatives to one another, and are often symbiotic in larger ecosystems. Many of the largest financial firms previously spent years questioning where this space would go, but now the numbers speak for themselves: 50% of the top 100 banks already have crypto and blockchain exposure, and nearly all plan to expand into digital assets. Just as with the Fintech explosion and the demand for digital experiences before it, crypto has the opportunity to be a defining point for those that choose to evolve rather than perish.

How Confluent Helps the DLT Space

The growing ancillary businesses for DLT and digital currency have aligned themselves with motions that are shared directly with Confluent: distributed data architectures, real-time operations 24/7/365, and cloud-first models.

This is no surprise, as Apache Kafka and the notion of blockchains share many similar traits: immutability, replication, distribution, and decoupling of applications and clients, to name but a few. This complementary relationship means that DLT and Confluent function particularly well together to extend the functionality of a given DLT through sidechain or offchain activities like analytics, integrations

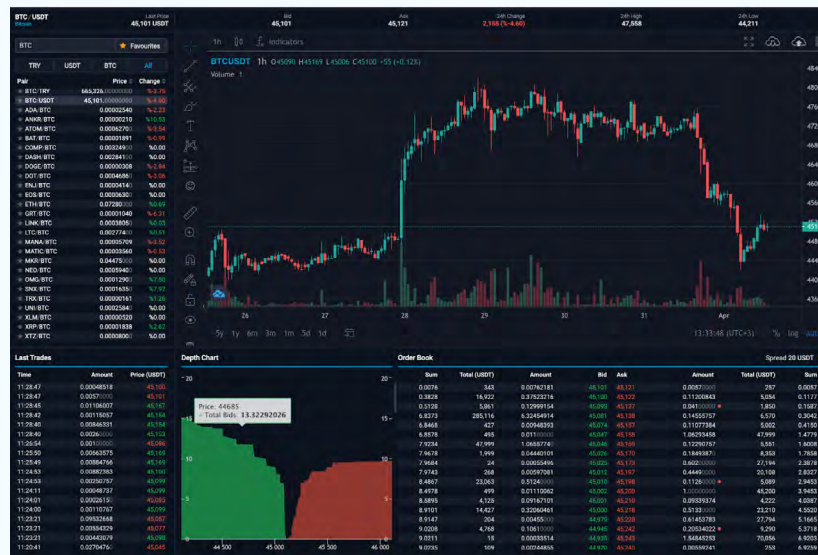
with traditional enterprise systems, or even integrating certain chains and ledgers together. With the blockchain technology space seeing a CAGR of nearly 50%, companies adopting these technologies require their data-in-motion platform to scale reliably with their own rapid growth—which is why they turn to Confluent.

Trading and Exchanges

A given block in the Bitcoin blockchain may contain data representing a transaction of the underlying asset, but there will also be a sale opposing this buy; a service, alternative cryptocurrency, or fiat currency transaction in exchange. Cryptocurrency exchanges must maintain a record of these orders as part of their orderbook. Suddenly the on-chain data is only a fraction of the data required for handling exchanges. Records of these trades naturally need to be stored with idempotency and strict ordering; that means only sent and processed once and stored in the order that they are received. This is important, as it ensures it's not possible to trade double the amount intended or allow users to spend beyond the limits of their reserves. As an order book is in essence a log, it has a natural compatibility with integration into topics that can be configured with the requirements mentioned above.

There are additional considerations that add complexity to handling crypto order books as compared to traditional trading markets. Compared to Foreign Exchange, cryptocurrency trading has much more variable demand and is more highly fragmented (EUR/USD makes up 23% of all FX transactions. In comparison, BTC/USD is at best 2% of crypto trading volume).

A familiar way to quickly interpret an orderbook is via a market-depth chart. The raw order data in the topic is too granular and lacks sufficient dimensionality to be visualized, so it makes most sense to aggregate it in a way more suitable for human cognition. Likewise, in a software-driven world, a trading bot needs to optimize its computation of the optimum price by discounting outliers. Both situations require real-time aggregations of the data. Stream processing opens up the possibility of handling state in the form of windowing and aggregation that can be represented visually in real time while still allowing access to the granular event that sits beneath.



Know Your Transaction

The outlook for custodial cryptocurrency businesses remains positive as they refine themselves over time, but they face increased stringency as they become more visible to regulators and try to onboard larger institutions. These businesses, especially exchanges, face major risks and challenges around security and compliance to extend their growth.

A technology called Know Your Customer (KYC), a process of verifying the identity of customers, is a response to one such challenge. It has allowed traditional financial firms to ensure their private customers really are who they claim to be through a background checking mechanism, but it has not always been enough and can suffer from silos of knowledge.

The exposure of transaction history allows a significantly deeper understanding of the provenance of a given fund transfer by exploring not just the user's verified profile but their current and historical transaction data, and comparing this to the baseline of their demographic, and equally exploring the networking leading to this given activity, regardless of the exchange or source from which it started.

KYT requires algorithms that can track on-chain activity to calculate risk levels based on associating suspicious activity, addresses, and wallets. It may additionally compare this against gambling services, money laundering systems, and so forth to score transactions and assess risk. If the assigned risk is too high, whoever is providing the virtual asset service is able to stop the commit or share the details with the individuals who have required an audit. Some KYT providers have even gone as far as to supplement this data with other external services (like social media) to perform activities like Natural Language Processing to measure the overall sentiment related to the given currency or protocol.

The ability to collect this data and process it requires a reliable high throughput technology that can allow transactions without delay in a highly volatile market. Kafka Streams, a library for building streaming applications, bypasses the need for computationally complex risk calculations at the point of request by instead building real-time fraud metrics available in a microservice with predictable scalability. This means that KYT processes can perform much more in real time.

TokenAnalyst is able to show this in a slightly higher-level view when their data science team uses ksqlDB to experiment with raw or lifted streams to ultimately deploy new machine-learning models (using custom user-defined functions) without writing a single line of Java code. Use cases range from stateless transformation, such as ERC-20 transactions, to anomaly detection and classification.

<https://www.confluent.io/en-gb/blog/reliable-fast-access-to-on-chain-data-insights/>

Integrating

For any technology to reach the point of true adoption, it must cross the chasm of being an isolated innovation product and become integrated into the wider ecosystem. Even in scenarios where distributed ledgers will eventually replace traditional systems, they will continue to need to be exposed to "the real world."

Perhaps the most mundane but important scenario of this would be integration with monitoring and observability tools. There is an operational cost to hosting a DLT node, so a way of sharing the data and state of the network to resources elsewhere in the business can be a force multiplier for the value generated. In this way, the experiences will be no different to traditional finance, where the private ledgers still require exposure to certain other systems, such as CRM.

Likewise, there are external data streams that can enrich the value of onchain data as both are fed into downstream systems. The smart contracts that are integral to the programmatic logic in a distributed ledger can be made into a hybrid by connecting non-blockchain data and systems into the on-chain process. Building around a smart contract requires it to be fully autonomous, which requires reliable exposure to real-world variables. These could be synchronous feeds, such as geo-location in relation to the purchase of a physical asset that is recorded to the chain, or asynchronous data, like a credit score that can be stored and then aligned at the point of a transaction to prove the validity of a transaction.

One category of digital assets most appropriate for this kind of setup is Defi, or Decentralized Finance. Over 200B\$ is locked up in distributed financial services—applications written into smart contracts that allow anyone in the world to lend, borrow, save, and trade digital assets without specific permissions. By doing away with a proverbial bank manager to validate the loan or set the interest rate, these applications require extra logic to protect the involved parties. This logic is represented in the form of oracles: middleware entities connected to the DeFi application's smart contract to provide off-chain data from outside their native blockchain. These could be liquidity aggregators, algorithmic stablecoins, or lending market data. These oracles themselves are capable of being distributed to protect the quality of the data, but they, too, need to be fed, and that is where Kafka can become a powerful source to feed into the likes of the API at DIA. Ultimately, event streaming becomes the enrichment middleware for all of the components in the distributed app.

<https://www.diadata.org/>

Confluent Outcomes in DLT

- Decreased risk of fraud—already a \$600M impact on the crypto market—through real-time analytics of DLT transaction data using event streaming.
- Reliably integrate off-chain and on-chain data sources to enrich the logic contained in smart contracts.
- Introduce DLT data to the wider organization through simple, scalable data streaming pipelines with 120+ pre-built connectors for real-time integration between source and destination system.
- Expose and monetise filtered DLT data to off-chain clients, with replayability, without having to re-download the entire chain history.
- Refine positions through aggregating trade histories in a way that is both more human-readable and better optimized for algorithmic trading.
- A flexible data platform adopted by traditional finance and decentralized finance businesses alike.

Glossary

Distributed Ledger Technology A database that is consensually shared and synchronized across multiple sites, institutions, or geographies, accessible by multiple parties to witness and own.

Blockchain An append-only store of data that records linked transactions distributed across a network. The linking is a control mechanism where the changes of each transaction are hashed to make sure the chain cannot be modified or added without being cryptographically validated. The blockchain is visible to all the peers in the network.

Distributed App Distributed applications are built on or around the smart contract in a blockchain to provide extra utility. This can be an interface to the blockchain data, and is the code running on the decentralized node in the peer-to-peer environment compared to the traditional notion of running an application locally. The internet of these apps is often referenced as or considered part of Web 3.

Non-Fungible Token (NFT) NFTs are tokens on a blockchain network used to represent ownership of unique items. Fungibility is the property of an asset or good to be interchangeable amongst its individual units.

An NFT can only have one owner at a time. Ownership is managed through a uniqueID and metadata that no other token can replicate. NFTs are minted through smart contracts that assign ownership and manage transferability.

Smart Contract Programs that are stored on a blockchain. Smart contracts define rules programmatically, like a regular contract, and automatically enforce them. The code and the action of the smart contract is distributed just as the data is, so they cannot be deleted by default and the interactions are irreversible. They remove the need for an intermediary and can greatly increase the speed of a transaction. They can also be a trigger for a workflow, kicking off additional actions when a condition is met.

Compound Annual Growth Rate (CAGR) Compound Annual Growth Rate, based on a formula that computes the average rate of growth between two given years. This allows the ability to compute growth rates more accurately in lieu of traditional arithmetic averages that will be more affected more by extreme swings during exponential growth.

Semantic Web/Web 3 The original definition of the semantic web by Tim Berners Lee: a “web of data that can be processed directly and indirectly by machines.” The interpretation of this from the perspective of DLT technology is reflected in “Web 3,” the decentralization of user-created data from specific organizations to wide availability for individuals and applications to interact with and build from.