THE ULTIMATE DATA STREAMING GUIDE

Concepts, Use Cases, Industry Stories

Guest chapter by Evi Schneider: How to build a data streaming community in your own organization

Kai Waehner
with a foreword by
Jay Kreps

CONFLUENT

THE ULTIMATE DATA STREAMING GUIDE

Concepts, Use Cases, Industry Stories

Kai Waehner
with a foreword by
Jay Kreps

Guest chapter by Evi Schneider: How to build a data streaming community in your own organization

Contents

Foreword by Jay Kreps 12

Preface 14

Content Matrix 18

01 Event-Driven Architecture and the Demand for Real-Time Data 20

1.1 The Business Impact of a Data Mess 22
1.2 The Need for an Event-Driven Architecture 24
1.3 Data Streaming Platform 26
1.4 Business Value of Data Streaming 30
1.5 Analytic Data Platforms 34
1.6 The Shift Left Architecture 36

02 Data Streaming Use Cases 38

2.1 Data Products 39
- Relation of Data Products to Data Streaming 39
- Examples for Data Products and Data Streaming 40
- Data Products Build with Data Streaming for Data Consistency and Continuous Processing 41

2.2 IT Modernization 42
- Relation of IT Modernization to Data Streaming 42
- Examples for IT Modernization with Data Streaming 42
- IT Modernization with Data Streaming for Seamless Integration to Ensure Data Accuracy 44

2.3 Observability 45
- Relation of Observability to Data Streaming 45
- Examples for Observability with Data Streaming 46
- Observability Through Data Streaming for Optimized Reliability and User Satisfaction 47

2.4 Hybrid and Multi-Cloud 48
- Relation of Hybrid and Multi-Cloud to Data Streaming 48
- Examples for Hybrid and Multi-Cloud with Data Streaming 49
- Hybrid and Multi-Cloud with Data Streaming for Global Availability and Scalability 51

2.5 Event-Driven Microservices 52
- Relation of Event-Driven Microservices and Data Streaming 52
- Examples for Event-Driven Microservices and Data Streaming 53
- Event-Driven Microservices Using Data Streaming for an Evolving Architecture and Technology-Independent Integration 54

2.6 Data Sharing **55**
- Relation of Data Sharing and Data Streaming 55
- Examples for Data Sharing 56
- Data Sharing Through Data 57

2.7 Artificial Intelligence (AI) **58**
- Relation of AI and Data Streaming 58
- Examples for AI Using Data Streaming 59
- Data Streaming for More Powerful and Accurate Results in AI Use Cases 60

2.8 Internet of Things (IoT) **62**
- Relation of IoT and Data Streaming 64
- Examples for Data Streaming in the IoT World 64
- IoT with Data Streaming to Optimize the OEE and Provide Innovative Consumer Experiences 65

2.9 Data Governance **66**
- Relation of Data Governance and Data Streaming 66
- Examples for Data Governance in Streaming Data Architectures 67
- Data Governance on Top of Data Streaming for Continuous Compliance and Data Sovereignty 69

2.10 ESG (Environmental, Social, and Governance) **69**
- Relation of ESG and Data Streaming 69
- Examples for ESG Use Cases with Data Streaming 71
- Improved ESG Compliance and Alignment Through Data Streaming 71

03 Industry Success Stories for Data Streaming **73**

3.1 Financial Services **75**
- Scenario 1: Legacy Offloading 76
- Scenario 2: Core Banking 77
- Scenario 3: Real-Time Payments 78
- Scenario 4: Fraud Detection 79
- Scenario 5: Regulatory Reporting 80
- Real-Life Examples (Erste Group Bank, KOR) 81
- Conclusion 86

3.2 Manufacturing and Automotive **88**
- Scenario 1: Condition Monitoring and Predictive Maintenance 90
- Scenario 2: Supply Chain 91
- Scenario 3: Track-and-Trace 92
- Real-Life Examples (Paul Tech AG, BMW Group) 93
- Conclusion 99

3.3 Retail and E-Commerce **100**
- Scenario 1: Inventory Management 101
- Scenario 2: Loyalty Platforms 102
- Scenario 3: Personalization 103
- Scenario 4: Live Commerce 104
- Real-Life Examples (Mirgros) 105
- Conclusion 108

3.4 Telco and Media **110**
- Scenario 1: Network Monitoring and Alerting 112
- Scenario 2: OTT Services 113
- Scenario 3: Advertisements 114
- Conclusion 115

3.5 Gaming **116**
- Scenario 1: Infrastructure Operations 118
- Scenario 2: Game Telemetry 119
- Scenario 3: Monetization Network 120
- Conclusion 121

3.6 Government and Public Sector **122**
- Scenario 1: Citizen Services 124
- Scenario 2: Smart City 125
- Scenario 3: National Security 126
- Conclusion 127

3.7 Energy and Utilities **128**
- Scenario 1: Smart Grid Monitoring 130
- Scenario 2: Industrial Control System 131
- Scenario 3: Energy Trading 132
- Scenario 4: Decarbonization 133
- Real-Life Examples (aedifion, AMPEERS ENERGY, Uniper) 134
- Conclusion 143

3.8 Transportation and Travel **144**
- Scenario 1: Transport Management 146
- Scenario 2: Ride-Hailing 147
- Scenario 3: Travel Booking 148
- Real-Life Examples (Austrian Post, LKW Walter, Lufthansa, Schiphol Group) 149
- Conclusion 159

3.9 Insurance **160**
- Scenario 1: Workflow Orchestration 162
- Scenario 2: Full-Stack Contract Handling 163
- Scenario 3: Micro Insurance 164
- Real-Life Examples (Generali Switzerland) 165
- Conclusion 169

3.10 Healthcare and Life Sciences **170**
- Scenario 1: Medical Diagnostics 172
- Scenario 2: Pharmaceutical R&D 173
- Scenario 3: Caregiving 174
- Conclusion 175

3.11 Technology **176**
- Scenario 1: Platforms and Mobile Apps 178
- Scenario 2: Infrastructure Monitoring 180
- Conclusion 181

04 Creation and Positioning of a Data Streaming Organization and Community **182**

Guest chapter by Evi Schneider

4.1 The Journey to a Data Streaming Organization 184
4.2 Key Drivers for a Community Within a Data Streaming Organization 187
4.3 Fundamental Tactics for Developing a Data Streaming Community 190
4.4 Strategies for Long-Term Engagement in a Data Streaming Community 193

05 Data Streaming Is a Journey: A Roadmap to Success **198**

A Testament to the Value of Data Streaming

Foreword by Jay Kreps

As the co-creator of Apache Kafka® and co-founder and CEO of Confluent, I've had the incredible opportunity to witness the evolution of data streaming—from its early days, where many viewed it as just a kind of scalable message broker, to its position today as a cornerstone of modern data architecture. Throughout this journey one of my favorite parts of my job is getting to meet with users and customers who are actually putting the technology into practice. For me it has been fascinating to watch how this technology found applications across virtually every industry, for applications well beyond anything we originally imagined.

This book, The Ultimate Data Streaming Guide—Concepts, Use Cases, Industry Stories, is a fantastic resource for anyone looking to dive deep into the world of data streaming. It captures not only the technology but also the stories behind its adoption and the value it delivers.

Apache Kafka, which started as a low-level infrastructure layer for data, has evolved into the basis for a rich and powerful ecosystem. Today, *over 150,000 organizations* rely on Kafka, a testament to its robustness and versatility.

It's exciting to see the growing number of industry analysts recognize data streaming as a distinct and vital software category in evaluations like the *Forrester Wave and IDC MarketScape*. The recognition of data streaming as a new software category by research analysts underscores its significance. Traditional software players and emerging startups alike are joining the data streaming game, further validating its importance. For example, Confluent has reached a billion-dollar revenue run rate within just a decade—clear evidence that there is something happening here.

> **"Apache Kafka, which started as a low-level infrastructure layer for data, has evolved into the basis for a rich and powerful ecosystem."**

While Kafka remains central to enterprise architectures that demand decoupling, scalability, real-time capabilities, or flexibility, we're seeing similar growth trajectories with Apache Flink in the realm of stream processing. Flink is rapidly becoming the go-to standard for processing streaming data, often paired with Kafka to power cuttingedge applications that harness the latest in AI and analytics to power real time use cases. Together, these technologies are redefining what's possible in the data ecosystem.

Learning how to apply streaming can be a challenge, though. Although the technology has become far simpler to program to, and cloud services have removed much of the operational burden, there is still a mindset shift in moving from batch to streaming architectures. That is where this book comes in. It dives into real-world use cases across industries, offering practical insights into what makes data streaming successful in practices in real applications. These applications are often harnessing data in new and creative ways, and powering some of the most essential applications at the heart of modern businesses.
By bringing together these examples I think Kai has created a fantastic resource to help inspire and inform the rest of us in how this technology can be applied out in the real world. Whether you're an experienced technologist or just beginning your journey into data streaming, you'll find valuable insights and examples to guide you.

Jay Kreps
CEO and Co-Founder of Confluent

KAI WAEHNER

Data Streaming is a Journey

Preface by Kai Waehner

In the fast-paced world of technology, real-time data beats slow data in most use cases across all industries.

Whether it's the financial services sector, manufacturing, or gaming, the ability to process and act on data as it happens is transforming the way businesses operate. Imagine a financial institution preventing fraudulent transactions the moment they occur, a manufacturing plant optimizing its production line in real time to improve the overall equipment effectiveness (OEE), or a gaming company delivering personalized experiences to players while they engage with the game. These are just a few examples of how data streaming is revolutionizing industries.

When you ask business leaders about their needs, without the constraints of specific technologies, the answer is clear: they want solutions that address their challenges effectively and efficiently. Data streaming and processing data in motion represent a paradigm shift from the traditional approach of storing data at rest and relying on databases and point-to-point APIs for request-response communication and delayed batch decisions.

This shift allows for the unification of operational and analytical workloads, all centered around an event-driven architecture.

At the core of this transformation are data products that are real time, of high quality, and scalable. Apache Kafka is used by over 100,000 organizations and has become the de facto standard for data streaming, while Apache Flink is rapidly growing to become the standard for stream processing.

These technologies complement existing data infrastructures like data lakes, data warehouses, and lakehouses. Apache Iceberg is emerging as the de facto standard for an open table format, enabling organizations to store data once and utilize their preferred analytics engine for each use case.

The business value of data streaming is undeniable. It is crucial to start your evaluation from the business perspective, focusing on solving real business problems. A purely technical proof of concept, as seen in the Hadoop era, is insufficient. Ideally, each use case and project should demonstrate tangible business value, such as increased revenue, reduced costs, minimized risks, and enhanced customer satisfaction.

One often underestimated advantage of data streaming is the true decoupling it offers, enabling a future-ready, evolving architecture and ensuring data consistency across real-time and non-real-time systems. This leads to faster time to market for IT modernization initiatives, such as mainframe offloading to the hybrid cloud, and innovation, like generative AI for new business models.

Data streaming is not just about open-source frameworks, commercial products, and cloud services; it also involves expertise, including support, 24/7 SLAs, and consulting. Evaluating different vendors for product features, SLAs, and other criteria is essential to finding the right fit for your organization.

> **"Data streaming and processing data in motion represents a paradigm shift."**

This book is structured to guide you through the world of data streaming:

- **Chapter 1:**
  Event-Driven Architecture and the Demand for Real-Time Data
- **Chapter 2:**
  Technical Data Streaming Use Cases
- **Chapter 3:**
  Industry Success Stories for Data Streaming
- **Chapter 4:**
  Creation of a Data Streaming Organization and Community
- **Chapter 5:**
  Data Streaming is a Journey: Conclusion and Outlook

The book includes a matrix of technical data streaming use cases and vertical industry success stories, completed with customer examples. This structure is designed for easy access, whether you are responsible for data streaming platforms and services within your company or need to present specific examples to upper management. IT executives can leverage the book's insights to highlight business value findings. For those reading the digital version, links are provided to seamlessly navigate from a case study in chapter 3 to the related technical use cases in chapter 2, ensuring a comprehensive understanding of how data streaming can drive success in your organization.

Data streaming is a journey, and by reading this book, you will learn about its transformative potential and how it can boost your business into the future while driving industry innovation. Through a maturity model, the book guides you from initial use cases and adoption across various lines of business, until data streaming, powered by an event-driven architecture, becomes the central nervous system of your organization.

Kai Waehner
Field CTO, Confluent

BETTER TOGETHER

Feedback? Questions? Get in Touch!

General Feedback or Questions:
use-case-book@confluent.io

Kai Waehner
Data Streaming Trends, Use Cases, Industry Stories:
LinkedIn, X (Twitter), Blog

Evi Schneider
Internal Community Creation, Positioning of Data Streaming Organization: LinkedIn

Content Matrix

Industry Success Stories for Data Streaming Data Streaming Use Cases	Financial Services	Manufacturing and Automotive	Retail and e-Commerce	Telco and Media	Gaming	Government and Public Sector	Energy and Utilities	Transportation and Travel	Insurance	Healthcare and Life Sciences	Technology
Data Products	Erste Group Bank		Migros					Austrian Post			
IT Modernization	Erste Group Bank KOR							Lufthansa Schiphol Group	Generali Switzerland		
Observability	KOR		Migros				aedifion	Austrian Post LKW Walter			
Hybrid and Multi-Cloud		BMW Group Paul Tech AG					AMPEERS ENERGY		Generali Switzerland		
Event-Driven Microservices			Migros				Uniper	Austrian Post LKW Walter			
Data Sharing	Erste Group Bank						Uniper	LKW Walter Lufthansa Schiphol Group			
Artificial Intelligence (AI)		BMW Group						Lufthansa			
Internet of Things (IoT)		BMW Group Paul Tech AG	Migros				AMPEERS ENERGY aedifion	Schiphol Group			
Data Governance	Erste Group Bank KOR								Generali Switzerland		
ESG (Environmental, Social, and Governance)		Paul Tech AG					aedifion	Schiphol Group			

To be extended!

—

If you want to become a part of it, don't forget to reach out to use-case-book@confluent.io

Event-Driven Architecture and the Demand for Real-Time Data

01

Frequent online shoppers are all too familiar with the following scenario: You find a product, place your order, and complete the payment. With every click, the anticipation of receiving your new pair of pants, your fancy espresso machine, or rowing ergometer grows—until you discover that the item is, unfortunately, no longer in stock because another espresso connoisseur or sportive buddy was quicker. It's frustrating. But is this just something we, as buyers, have to get used to because it's unavoidable? Are there simply "natural limits" to the availability of information that we must accept? Or are we—as private consumers, and even more so as business users handling data every day—on the threshold of a paradigm shift in how we use data?

Let's be clear right from the start: we absolutely don't have to accept delayed data as our fate. In fact, it could almost be stated as a universal rule: real-time data is better than slow or delayed data in nearly every application.

> **"Real-time data is better than slow or delayed data in nearly every application."**

Accurate and up-to-date information empowers organizations to make informed decisions, enhance customer experiences, and drive innovation. However, unlocking data and analytics is much harder than it should be. In the current digital landscape, most enterprises across all industries are grappling with the limitations of fractured and siloed data architectures. These traditional systems, characterized by their reliance on traditional databases, periodic data processing, and point-to-point APIs, often fall short in meeting the demands of modern business operations. The inherent latency and complex integration across systems means that data is only updated at scheduled intervals, leading to significant delays in decision-making and operational responsiveness.

This lag can be particularly harmful in use cases where real-time data is crucial, such as in the financial sector, where the inability to access real-time data can hinder risk management and trading strategies, while in healthcare, delayed data can impact patient care and treatment outcomes.

1.1 The Business Impact of a Data Mess

One of the primary challenges of a siloed data world is the inconsistency of information. As data is collected and processed in batches or through APIs, there is always a time gap between the occurrence of an event and its reflection in the system. This delay can result in outdated or stale data being used for critical business decisions and lead to suboptimal outcomes.

Most enterprise architectures store data across two distinct and different estates:

1. The **operational estate** which serves the transactional and real-time systems that run the business.
2. The **analytical estate** for after-the-fact analysis and business decision-making.

Organizations typically employ a mish-mash of different approaches for data integration to connect the operational and analytical estates. Moreover, batch-driven architectures are often built on monolithic legacy systems that are inflexible and difficult to scale. These systems are typically tightly coupled. Any change in one component can have cascading effects on the entire system. This rigidity makes it challenging to adapt to new business requirements or integrate with modern technologies. For example, in the fast-paced world of e-commerce, the inability to quickly adapt to changing customer preferences or market trends can result in lost sales and decreased customer satisfaction.

The "data mess" is further made worse by the increasing volume and velocity of data generated by modern enterprises. Traditional systems struggle to keep up with the sheer scale of data, leading to performance bottlenecks and inefficiencies. In industries like manufacturing, where the Internet of Things (IoT) generates vast amounts of data from sensors and equipment, the inability to process this data in real time can disrupt typical use cases like predictive maintenance, ultimately hindering operational efficiency. This complexity leads to a so-called "spaghetti architecture."

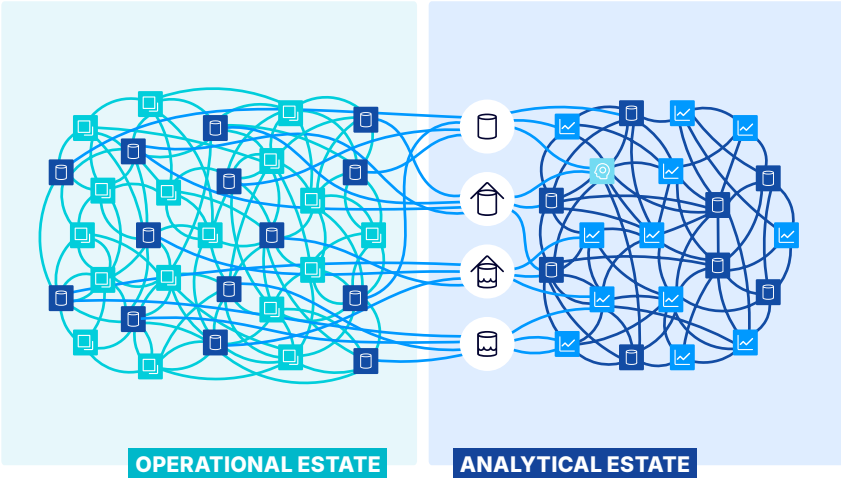

Spaghetti architecture: Data complexity escalates into a data mess as data needs multiply over time

In contrast, data streaming offers a transformative solution to these challenges. By enabling continuous data processing and immediate insights, data streaming ensures that businesses can respond rapidly to emerging technical or business events. This capability is invaluable across all industries and provides a competitive edge in a data-driven world.

The shift from batch systems, databases, and APIs to data streaming is not just a technological upgrade but a strategic imperative to solve the giant data mess. Embracing data streaming can significantly enhance operational efficiency, reduce costs, and unlock new revenue streams. By investing in modern data architectures, IT executives can ensure their organizations remain agile, competitive, and positioned for future growth.

"By investing in modern data architectures, IT executives can ensure their organizations remain agile, competitive, and positioned for future growth."

1.2 The Need for an Event-Driven Architecture

IT architectures are currently undergoing a significant evolution: the shift from tightly coupled point-to-point communication using databases and APIs to event-driven architectures (EDA). This marks a major change in how modern organizations design and manage their systems. While traditional point-to-point integrations may serve their purpose, they often lead to a web of dependencies which is difficult to manage and scale. In such systems, each application is directly connected to the others, creating a tightly coupled structure where changes to one application can affect the entire enterprise architecture. This inflexibility hinders innovation and makes it challenging to adapt quickly to new business needs and achieve rapid time to market.

In contrast, an event-driven architecture offers true decoupling of applications, enabling a more flexible and scalable system design. In an event-driven architecture, applications communicate through events, which are discrete pieces of information that signify a change in state. These events are published to a central event bus where they can be consumed by any interested consuming application.

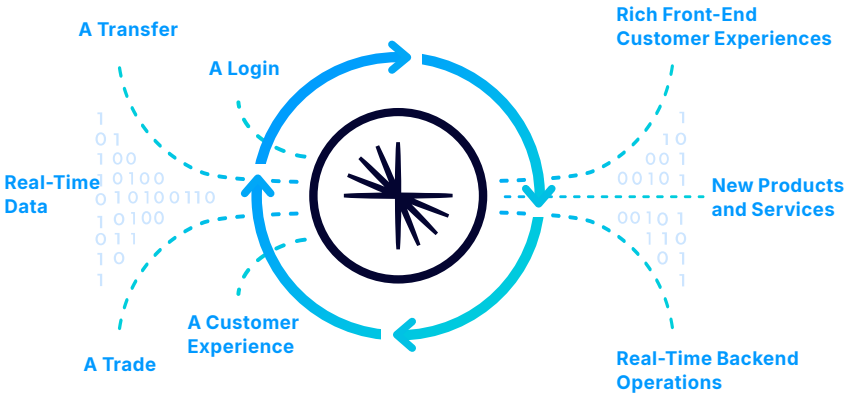

The need for an event-driven architecture to set data in motion

This decoupling using an event-driven architecture allows each application to:

- **operate independently**
- **reduce dependencies**
- choose the **right technology**
- **make it easier and faster** to introduce new features or services

One of the key benefits of an event-driven architecture is its alignment with domain-driven design (DDD). DDD is a software development approach that focuses on modeling software to match the business domain. By using events to represent business processes and changes, an EDA naturally supports the principles of DDD, enabling a more intuitive and maintainable system design. This alignment ensures that the technical architecture closely mirrors the business processes, facilitating better communication and collaboration between technical and business stakeholders.

Furthermore, an event-driven architecture is a key component of a data mesh, a modern approach to data management that treats data as a product. A data product is a curated asset designed to deliver specific value or solve particular problems for stakeholders. Unlike raw data, a data product is structured, refined, and packaged in a way that it can be easily consumed, interpreted, and utilized by its intended audience. In a data mesh, data ownership is decentralized, with each domain responsible for its own data products. Events play a crucial role in this model, as they serve as triggers of data sharing between the decoupled systems. Events are pre-defined changes in state and can be anything from a product being ordered, a user address being edited to other stimuli (like an alert issued by security systems).

"Events play a crucial role, as they serve as triggers of data sharing between the decoupled systems."

In an EDA, events allow data products to be immediately available in the connected systems, making data sharing and integration across domains more seamless. This approach not only improves data quality and accessibility but also fosters a culture of data-driven decision-making within the organization.

1.3 Data Streaming Platform

Data streaming is a fundamentally different paradigm, elevating data to a first-class citizen. Unlike traditional batch processing or point-to-point communication via APIs, data is not static or passive. It's constantly evolving and continuously being shared in real time. The fundamental notion is that data is moving, not sitting at rest waiting to be queried or to be extracted.

Data streaming enables the continuous flow of data throughout a business. This capability is particularly valuable in today's fast-paced business environment, where timely information can be a critical differentiator.

How does data streaming help solve the giant data mess and get to a central nervous system?

Data Streaming enables teams to package and serve their data as if it were a ready-to-consume product, **one that can be put to immediate use.**

1. It shifts your thinking from tightly coupled point-to-point integrations to a **multi-subscriber**, decoupled architecture leveraging data products.
2. Instead of repeatedly wrangling and cleansing data downstream, data quality and curation are shifted upstream. This allows domain experts **to present ready-to-consume data that is governed, trustworthy, and reusable**, enabling consumers to instantly apply those datasets to their use cases.
3. Instead of batch-based integration models like ETL or ELT, business units focus on **a continuous flow of integrated data** that is always up to date.
4. And finally, instead of querying your data that's **fixed in time, processing, analyzing, and acting** on data happens continuously.

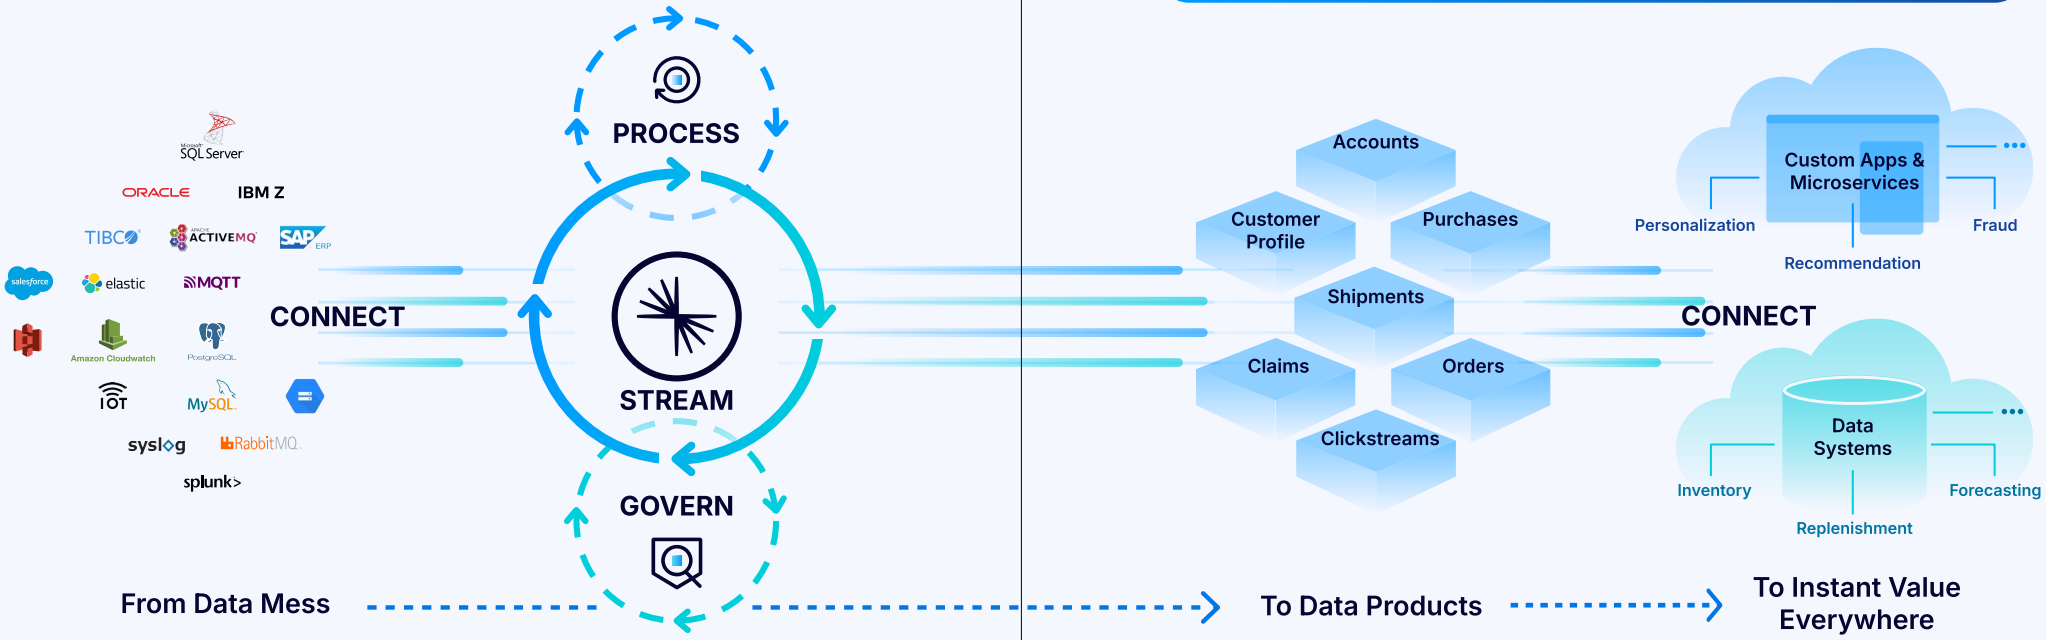

A data streaming platform powered by an event-driven architecture to enable reusable, real-time, and trustworthy universal data products

This approach enables teams to package and serve their data as if it were a ready-to-consume product, one that can be put to immediate use.

Hence companies rely on data streaming as the foundation for their data product strategies. Data products should be reusable assets that represent key business entities—such as customers, accounts, claims, inventory, and shipments. And when done right, they are constantly up to date (real time) and continuously governed as well as reusable across both operational and analytical estates. Most data streaming infrastructures are powered by the open-source technologies Apache Kafka and Apache Flink.

Apache Kafka is a distributed data streaming platform that enables the real-time collection, storage, and processing of data. It is widely adopted across industries.

Over 100,000 organizations are leveraging Kafka's capabilities to build scalable and resilient data pipelines. The robust architecture and high throughput make it an ideal choice for both handling large volumes of data and critical transactional workloads with low latency. The persistence layer enables data consistency across real-time and non-real-time systems such as databases, data lakes, file systems, and request-response APIs.

Apache Flink, on the other hand, is a stream processing framework that complements Kafka by providing powerful capabilities for continuous processing of operational and analytical workloads. Flink's advanced features, such as stateful processing and complex event processing (CEP) enable real-time analytics at extreme scale. Together, Kafka and Flink form a powerful combination for building end-to-end data streaming solutions.

While Kafka and Flink provide the foundational technologies for data streaming, Confluent takes it a step further by offering a complete data streaming platform. Confluent's platform extends the capabilities of Kafka and Flink with additional features and tools that simplify the development, deployment, management, and operations of data streaming pipelines and applications.

Confluent Cloud, a fully managed service, provides a cloud-native platform that runs on all major cloud providers. This flexibility allows organizations to leverage the power of data streaming without the operational overhead of managing infrastructure.

"Confluent Cloud, a fully managed service, provides a cloud-native platform that runs on all major cloud providers."

For organizations with the requirement to self-manage (usually because of cost, security, or latency reasons), Confluent Platform offers a comprehensive solution that can be deployed in data centers and public cloud customer VPCs, or even at the edge outside of a data center, such as in factories or retail stores. This versatility ensures that businesses can harness the benefits of data streaming regardless of their deployment environment.

As a third option in the Confluent product portfolio, WarpStream offers a cloud service utilizing the Kafka protocol and provides organizations with the flexibility of a Bring Your Own Cloud (BYOC) model. This approach allows businesses to maintain control over their data by deploying WarpStream in their own cloud environments while benefiting from the managed service's operational support and scalability. By combining the advantages of a managed service with the security and customization of a self-hosted solution, WarpStream's BYOC model empowers organizations to optimize their data streaming strategies according to their unique security and compliance requirements.

The adoption of a data streaming platform represents a strategic investment in the future of their organization's data architecture. By enabling real-time data products, companies are equipped to increase operational efficiency, enhance customer experiences, and unlock new revenue streams. The ability to deploy data streaming solutions across multi-cloud, on-premise, and edge environments ensures that organizations can remain agile and responsive to changing business needs and compliance requirements. Streaming and processing events is the core of a data streaming platform. Additionally, connectors play a crucial role, serving as the bridge between various data sources and sinks. They facilitate the seamless integration of disparate systems, enabling the flow of data across different environments.

Connectors provide connectivity to legacy systems like a mainframe but also the latest technologies and cloud services like a cloud lakehouse or GenAI SaaS service. If no connector exists, several options are available: Integration via APIs, clients for all major programming languages, implementing a custom connector, or leveraging a connector developed by a third-party vendor. Data governance is a critical add-on in a data streaming platform to ensure data quality, compliance, and security. It provides real-time data lineage and traceability and a metadata catalog for self-service consumption of existing data products to enable immediate actionability and consistent policy enforcement across distributed environments. By standardizing data formats and protocols, data governance facilitates seamless integration and supports strategic business initiatives, fostering trust and collaboration among stakeholders that use the same data product in different business units or projects.

1.4 Business Value of Data Streaming

In the realm of data streaming, technological capabilities are only part of the equation. Data products can be tapped into and joined to rapidly implement real-time capabilities that make big business impacts. But to truly harness the power of data streaming, organizations must understand and articulate the business value it brings. Even the most advanced technology will fall short if it does not align with business objectives and deliver tangible benefits. Therefore, it is crucial to focus on the business solutions and benefits that data streaming can underpin.

One of the primary ways data streaming can drive business value is by supporting new solutions that drive revenue. For example, in the retail industry, real-time data streaming can power personalized marketing and dynamic pricing strategies. By analyzing customer behavior and market trends in real-time, retailers can tailor their offerings to individual customers, increasing sales and customer loyalty. Similarly, in the financial sector, real-time data streaming can enhance trading strategies and risk management, leading to higher returns on investment.

Data streaming also plays a significant role in reducing costs. In manufacturing, for instance, real-time monitoring of equipment and processes can power predictive maintenance, reducing downtime and maintenance costs. By identifying potential issues before they lead to equipment failure, manufacturers can optimize their operations and extend the lifespan of their assets. In the healthcare industry, real-time data streaming can support the streamlining of administrative processes and improve patient care, leading to cost savings and better health outcomes.

Overview on business solutions and benefits that data streaming can underpin

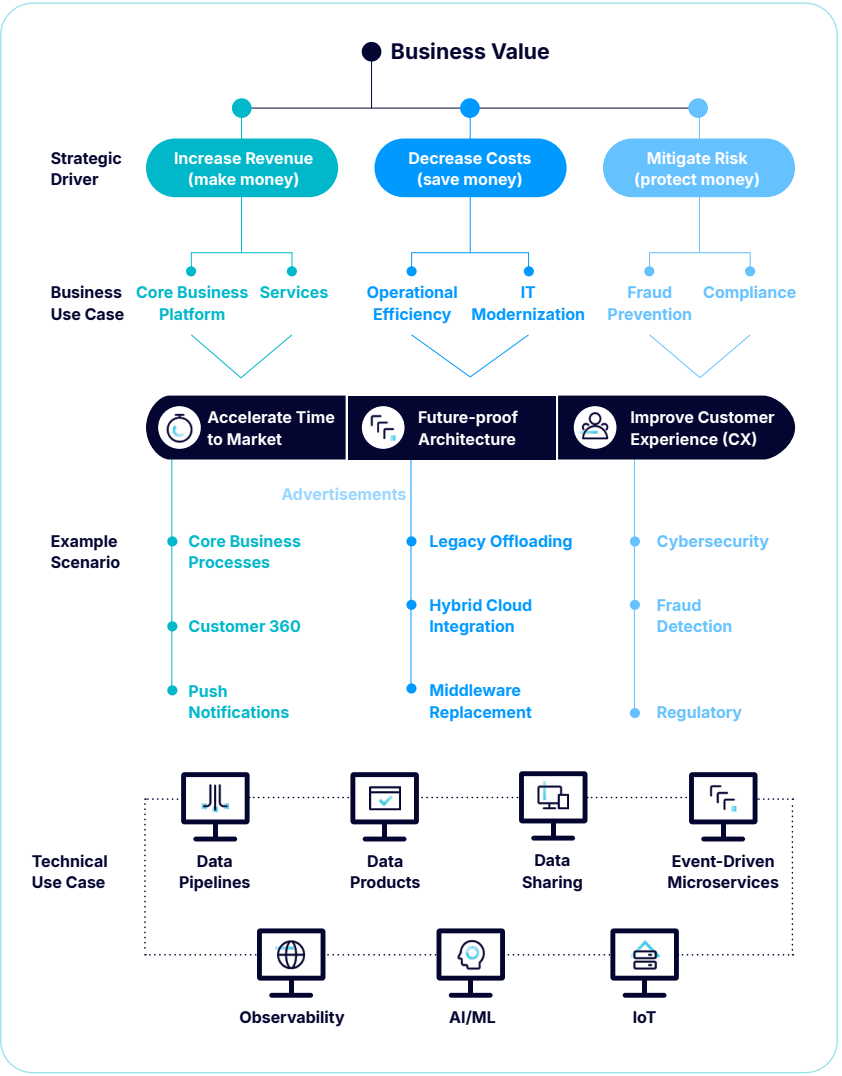

Strategic drivers for data streaming initiatives and related use cases and scenarios.

Another critical aspect of business value is risk reduction. In the context of cybersecurity, real-time data streaming is required to provide immediate detection and response to threats, minimizing the impact of security breaches. By continuously monitoring network activity and analyzing data in real time, organizations can identify and mitigate risks before they escalate. In the financial industry, real-time fraud detection systems powered by data streaming can prevent fraudulent transactions and protect customer assets.

Understanding the business value of data streaming is essential for securing buy-in from stakeholders and driving successful implementations. By demonstrating how data streaming contributes to increasing revenue, reducing costs, and mitigating risks, IT executives and their teams can build a compelling case for investment in modern data architectures. This alignment between technology and business objectives ensures that data streaming initiatives deliver meaningful and measurable outcomes.

How much do you agree or disagree with these statements?

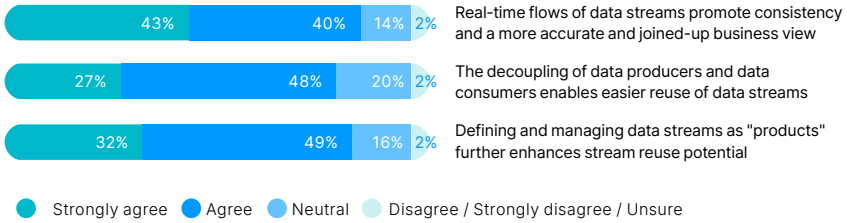

Insights from the Confluent Data Streaming Report 2024

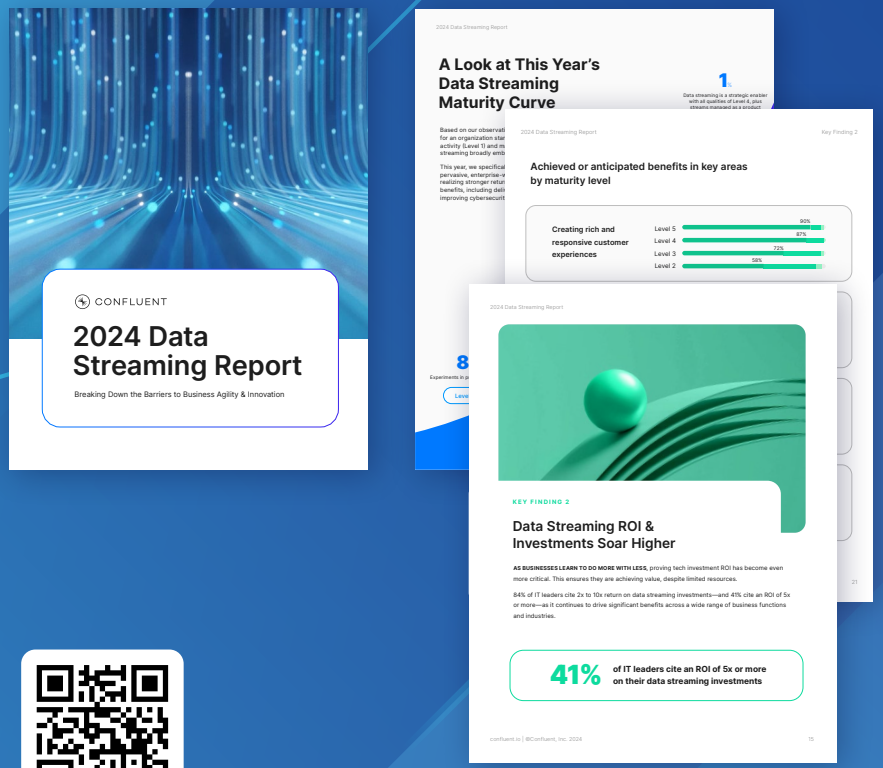

The *Data Streaming Report 2024* shows how IT leaders confirm the business value of data streaming in their organizations:

1 Event-Driven Architecture and the Demand for Real-Time Data

1.5 Analytic Data Platforms

The landscape of data platforms has evolved significantly over the years, with each type of platform offering unique capabilities and benefits.

Traditional data platforms primarily focus on building data pipelines that store and process data at rest. Data pipelines are the series of processes that move data from one system to another. The data is moved from sources (such as databases, IoT devices, or APIs), through an ETL (extract, transform, load) process where the data is extracted, cleaned, transformed, and enriched, before delivering it to data sinks (such as critical applications or analytical data platforms). Alternatively, ELT (Extract, Load, Transform) first extracts and loads data into a target system before it is transformed within that system.

Among those traditional data platforms are data warehouses and data lakes. Data warehouses are designed for structured data and are optimized for complex queries and reporting. They are commonly used for business intelligence and analytics, providing historical insights based on large volumes of data.

Data lakes, on the other hand, are designed to handle both structured and unstructured data. They offer greater flexibility in terms of data storage and processing, making them suitable for big data analytics and machine learning applications. However, data lakes can become unwieldy and difficult to manage if not properly governed. This leads to issues such as data swamps.

The Conventional Extract, Load, Transform (ELT) Architecture with Batch Workloads, Iconsistent Data, and High Compute Cost

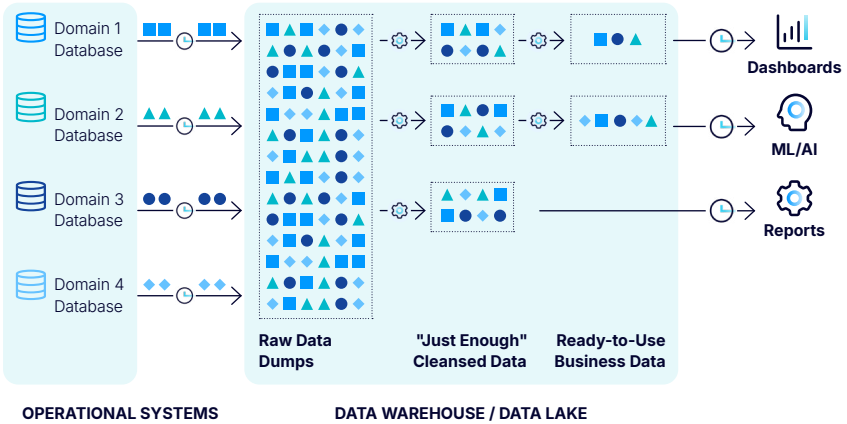

A conventional ELT architecture using a "lakehouse"

With the "lakehouse," a hybrid concept has emerged that aims to combine the best of both data warehouses and data lakes. A lakehouse provides a unified platform to build data pipelines for all types of data to enable both analytical and transactional workloads. This approach simplifies data management and improves data accessibility, making it easier for organizations to derive insights from their data. The conventional data processing of data lakes and warehouses often leads to batch workloads, inconsistent data and high compute costs because of repeated data processing.

While these traditional data platforms focus on data at rest for analytical use cases, data streaming introduces the capability to process data in motion. The widely adopted approach of feeding data from the lakehouse back into operational systems using Reverse ETL is considered an anti-pattern. It complicates data architecture by moving data back from analytical systems to operational systems, leading to increased complexity, potential data inconsistencies, and higher maintenance costs.

A data streaming platform enables continuous data pipelines to provide immediate insights and enable timely decision-making for operational and analytical applications and platforms. This capability is particularly valuable in scenarios where the speed of data processing is critical, such as real-time fraud detection, predictive maintenance, and personalized marketing. Additionally, a data streaming platform provides the required SLAs for uptime, availability, and transactional guarantees to support both analytical and operational workloads.

Understanding the interplay between different data platforms is crucial for designing a robust, scalable, and cost-efficient data architecture. By combining the capabilities of data warehouses, data lakes, lakehouses, and data streaming, enterprises can create a unified data ecosystem that supports both analytical and operational workloads. This holistic approach ensures that organizations can derive maximum value from their data, driving innovation and competitive advantage.

Data Streaming vs. Lakehouse

Data streaming and other data platforms like modern lakehouses, are complementary rather than competitive. Data products provide an interface to connect, process, and share data with any platform. Data streaming can enhance the capabilities of data warehouses, data lakes, and lakehouses by providing real-time data ingestion and processing. For example, data streaming can be used to continuously feed the data product into a data lake or data warehouse to ensure that the data is always up to date, consistent, and ready for analysis. Additionally, the data streaming platform connects not just to analytical data streams, but also the most critical transactional data set for processing and for ingesting into the lakehouse for analytics. This integration enables organizations to leverage the strengths of both real-time and batch processing, creating a more comprehensive and agile data architecture.

1.6 The Shift Left Architecture

The "Shift Left Architecture" is a modern approach to data management that decouples data storage from data processing and consumption. The approach emphasizes the early integration of data processing and data quality checks in the data pipeline. Following this idea, organizations conduct data processing and validation much sooner within the timeline ("shifting" them "to the left") than they would in other architectures—allowing them to identify and address issues earlier, in turn improving data quality and reducing cost by avoiding re-processing information multiple times and gaining faster access to insights and actions.

This concept enables greater flexibility and scalability. It allows organizations to choose the best tools and technologies for their specific needs.

Data streaming is a key enabler of the shift left paradigm, as it allows for continuous data processing and real-time validation.

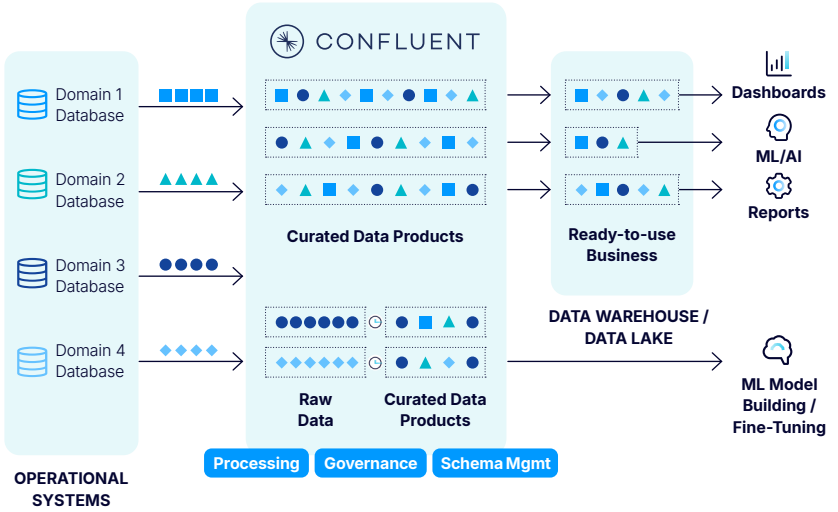

The "Shift Left Architecture" for real-time data, better data quality, and reduced compute cost

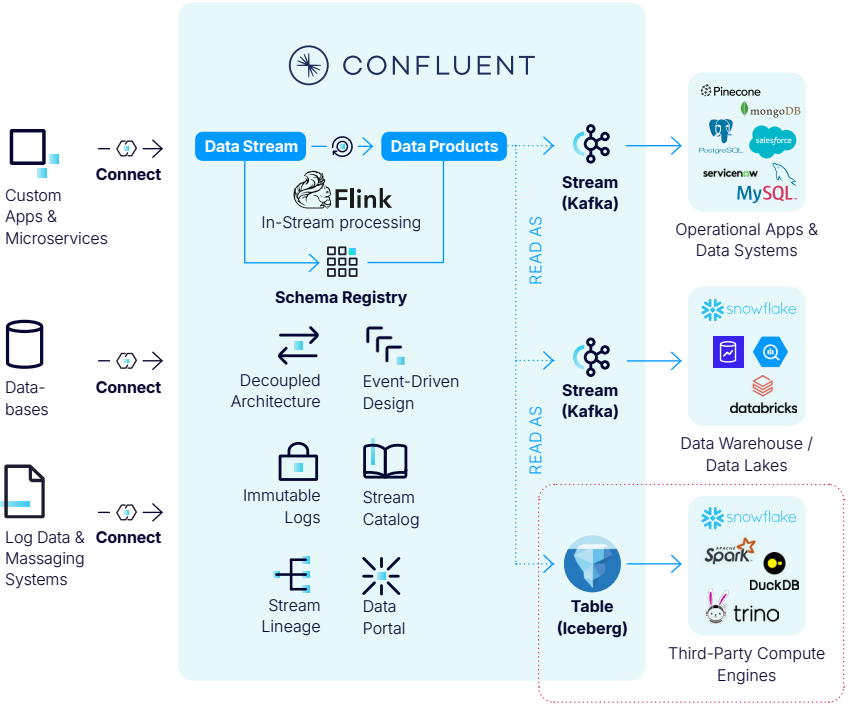

Concrete example of a Shift Left Architecture with Apache Kafka, Apache Flink, and Iceberg to unify operational and analytical workloads

Combining data streaming with the shift left architecture makes universal data products possible. These are reusable data assets that can be consumed by different operational and analytical applications and services across the organization in an event-driven architecture. By leveraging data streaming, organizations can create real-time data products that provide consistent and up-to-date information to enable better decision-making and operational efficiency.

Data products are the foundation for building new data pipelines in good data quality with any data platform. Apache Iceberg, the de facto standard for an open table format for large analytic datasets, plays a crucial role in this type of architecture by providing a consistent and efficient way to consume data across different compute engines from a single storage system.

Data Streaming Use Cases

02

Each organization begins its data streaming journey in its own way. The initial use cases for adopting data streaming are as varied as the industries they appear in. However, the core technical use cases often remain the same, showing up differently across business contexts depending on the industry. This chapter highlights the most common horizontal technical use cases seen over recent years, relevant to both those just starting and those already progressing in their data streaming adoption.

The following 10 use cases do not cover all of data streaming's technical capabilities worldwide, but they aim to provide a clear overview of the options available for starting and expanding a data streaming platform.

2.1 Data Products

While the term "data product" may appear too technical, perhaps even abstract, we find ourselves surrounded by them constantly. Whether we're grumbling about the recommendation engine for our product portfolio, proudly showing off our health data on our smartwatch or using a GPS system as decision support on our daily commute—data products are an essential cornerstone for a multitude of applications that we use in our daily routines.

Data products can take many forms, but at their core they are curated, reusable datasets created from data pipelines that provide specific, valuable insights or functionality to applications. Data products deliver targeted, actionable insights that allow data-driven integration, decision-making, and innovation, while data mesh enhances business agility and scalability within the organization. Many organizations are espousing the benefit of data products.

According to *McKinsey* and *Harvard Business Review*, the benefits of the data product approach can be significant:

- **New business use cases** can be delivered as much as 90 % faster.
- **The total cost of ownership**, including technology, development, and maintenance, can decline by 30 %.
- **The risk and data governance** burden can be reduced.

Relation of Data Products to Data Streaming

In an event-driven architecture, data is constantly ingested, processed, and delivered across the data pipeline to downstream consumers to ensure that information is always up to date. This real-time capability is essential for

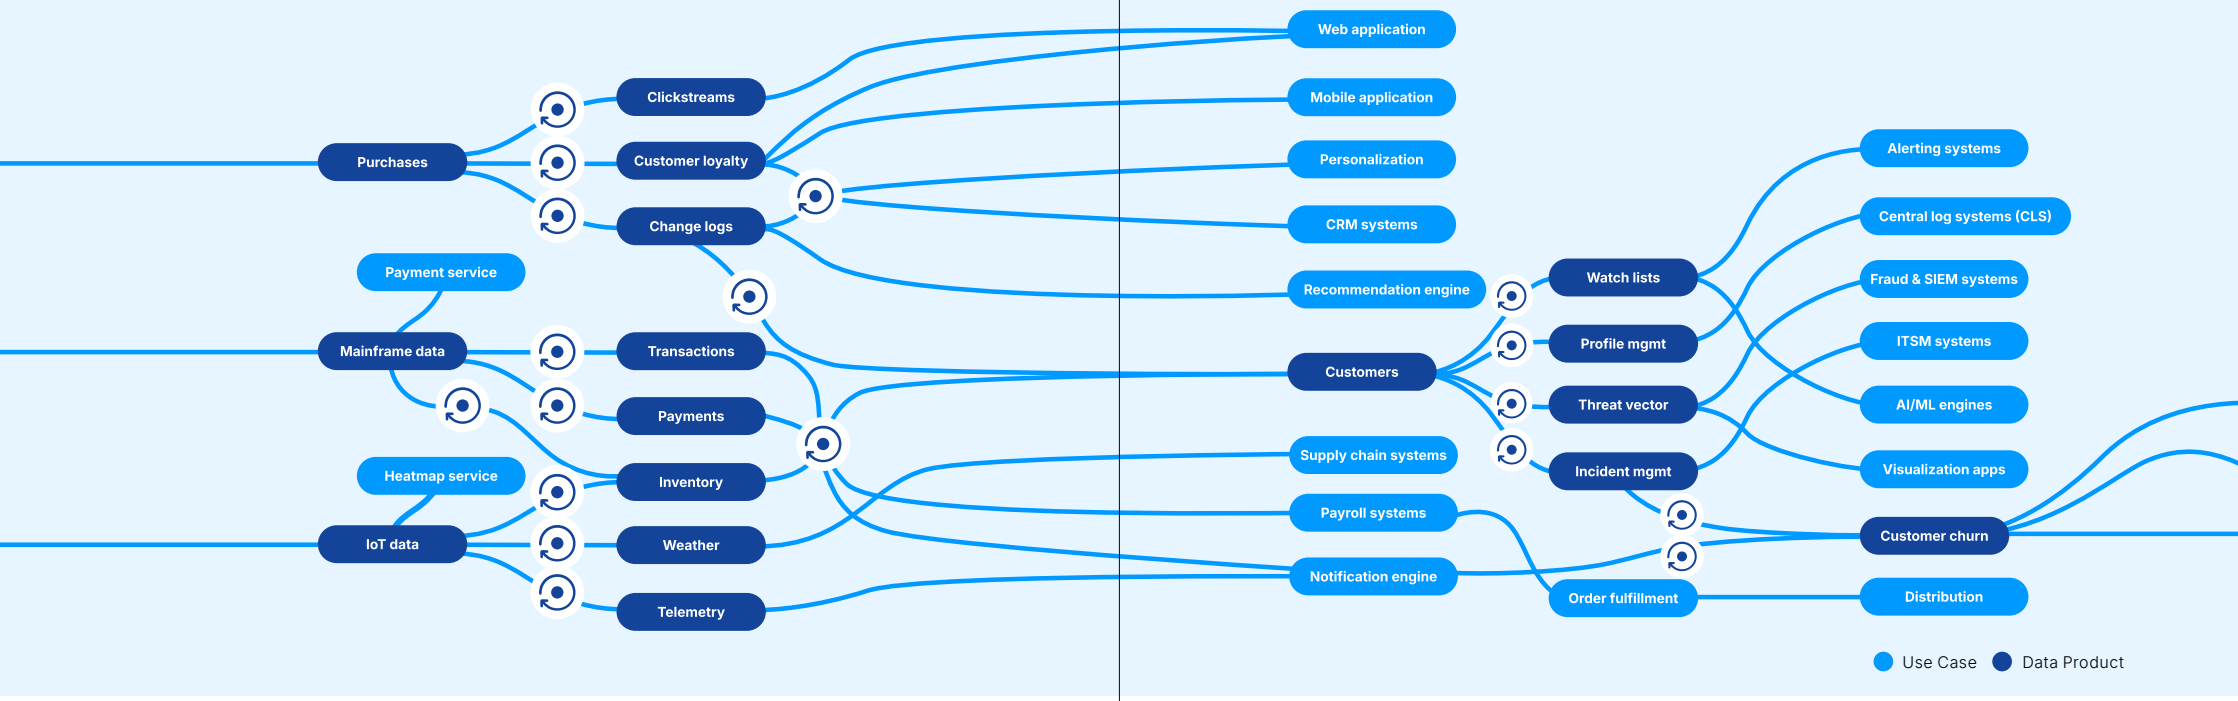

Reusage of data products across multiple use cases

maintaining high data quality and consistency across various applications, whether the producers or/and consumers are real time, batch, or API-based. Data streaming improves the quality of data for all downstream consumers by providing a consistent and reliable flow of information. This is achieved through the continuous ingestion and processing of data, which ensures that any changes or updates are immediately reflected across the entire data pipeline. By integrating data sources and sinks in real time, the data streaming platform enables seamless data flow and reduces the risk of data inconsistencies.

Examples for Data Products and Data Streaming

Imagine a retail company that uses data streaming to manage its inventory and sales data. By leveraging data streaming, the company can continuously ingest data from various sources, such as point-of-sale systems, online stores, and supplier databases. This real-time data is then processed and enriched to create data products, such as inventory reports and sales forecasts, which are consumed by different departments within the organization.

Another example is a financial institution that uses data streaming to monitor transactions and detect fraudulent activities. By integrating data from multiple sources, such as transaction logs, customer profiles, and risk models the institution can create a comprehensive view of each transaction in a data product. This real-time data is then used to identify suspicious patterns and trigger alerts, enabling the institution to respond quickly and mitigate risks.

In a long-term data strategy, data products can and should be reused across use cases to increase time to market and reduce costs.

Data Products Build with Data Streaming for Data Consistency and Continuous Processing

Since data products are usually consumed by multiple downstream applications, data quality and governance are even more critical for compliance and observability reasons. Depending on their use case, they also need to be kept up-to-date all the time. Think of your music streaming app: you would not want it to keep recommending five more songs of the band you just marked as "Not interested", would you?

All these requirements are best met by adopting an event-driven architecture along with data streaming to ensure that data products are always accurate, consistent, and readily available to all business units. That way, organizations can make the most out of their data by enhancing data quality, improving operational efficiency, and enabling real-time decision-making across the data pipeline.

2.2 IT Modernization

IT Modernization involves updating and transforming an organization's IT infrastructure, applications, and processes to leverage modern technologies and practices.

This process often includes migrating legacy systems to the cloud, adopting microservices, and implementing advanced data architectures. Data integration plays a crucial role in IT modernization by enabling seamless connectivity and data consistency between old and new systems to ensure that data flows smoothly across various platforms uni- or bidirectional.

Data consistency is essential in this context, as it ensures that data remains accurate, reliable, and synchronized across the modernized IT landscape to support informed decision-making and operational efficiency no matter what technology, programming language, API, or communication paradigm each system uses.

Relation of IT Modernization to Data Streaming

By providing an event-driven architecture with real-time data integration and processing capabilities, data streaming supports the transition from legacy systems to modern, event-driven applications. This shift is crucial to maintain data consistency and ensure that information is always up to date and accurate. In a hybrid cloud environment, data streaming platforms provide seamless integration and synchronization between on-premise systems and cloud-native services. This capability is particularly important for organizations that need to modernize their IT infrastructure while maintaining compatibility with legacy technologies, such as mainframes, and integrating with modern applications, both self-developed and commercial products or SaaS.

Examples for IT Modernization with Data Streaming

Imagine a healthcare organization that is modernizing its IT infrastructure to improve patient care and operational efficiency. By leveraging data streaming, the organization can integrate data from various sources, such as real-time electronic health records (EHRs), near real-time high-volume logs from medical devices, and transactional data from external databases in batch. The event-driven data integration ensures that healthcare providers have access to consistent and up-to-date information, enabling better decision-making and improved patient outcomes.

Another example is a financial services company that is transitioning from a legacy mainframe system and monolithic middleware to decoupled, cloud-native applications. By using a data streaming platform, the company can continuously ingest and process data from its mainframe, ensuring that the information is always consistent and up to date via change data capture (CDC). This real-time data integration enables the company to leverage advanced analytics and machine learning capabilities to drive innovation and competitive advantage in an elastic cloud environment leveraging fully managed SaaS.

An additional example is the modernization of Security Information and Event Management (SIEM) systems. Traditional SIEM solutions often struggle with increasing volume and velocity of security log events. By leveraging data streaming, organizations can complement their SIEM solutions to preprocess data continuously and only feed relevant information into the SIEM to handle the backpressure to only ingest the supported throughput, and as an economic side effect reduce the overall SIEM cost significantly.

Data streaming can ingest and process security logs and events from various sources, providing a unified and comprehensive view of the organization's security posture to complement existing SIEM solutions. Some special SIEM use cases like situational awareness and threat detection require the stateful analysis of security events in real time at low latency to improve response times and prevent attacks before they happen. This is where most SIEMs are too slow and cannot handle the scale. Stream processing as part of the data streaming platform can be used to implement these use cases in separated applications, e.g., a Flink application.

IT Modernization:
Lift and Shift with Data Streaming

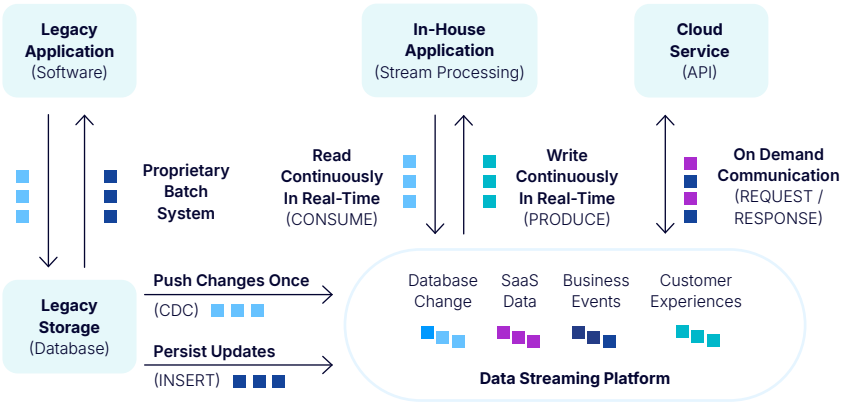

IT Modernization through an event-driven architecture with decoupled applications and a variety of integration options

IT Modernization Example: Mainframe Offloading and Replacement with Data Streaming

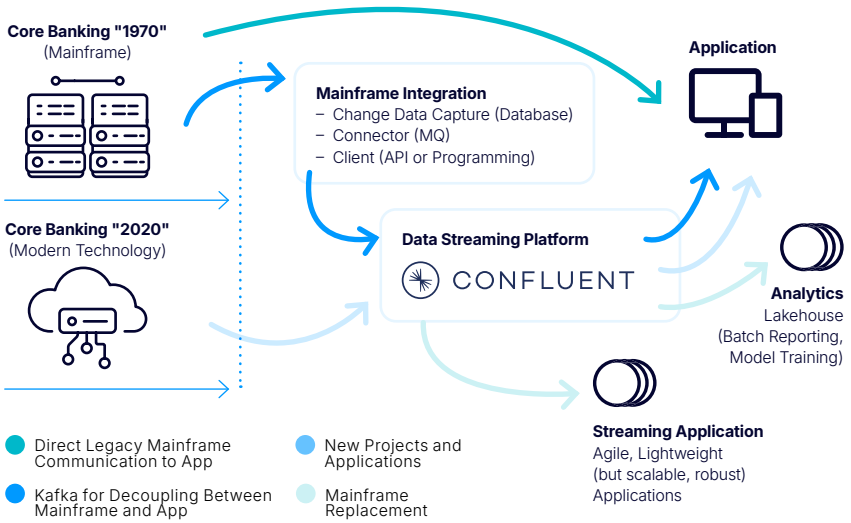

IT Modernization Example: Mainframe Offloading and Replacement with Data Streaming

IT Modernization with Data Streaming for Seamless Integration to Ensure Data Accuracy

The modernization of IT infrastructure through data streaming based on event-driven architectures offers significant business value by improving the scalability, flexibility, and data consistency of data pipelines. By leveraging data streaming, organizations can seamlessly integrate and synchronize legacy systems with modern cloud-native services to ensure that data is always accurate and reliable.

The adoption of a data streaming platform represents a strategic investment in the future of the organization's IT infrastructure. Adopting a holistic approach to IT modernization through data streaming significantly enhances operational efficiency and reduces costs by enabling true decoupling of systems, which allows for seamless integration of legacy and modern technologies. This flexibility empowers each business unit to adopt and modernize at its own pace, ensuring that technological advancements align with specific business needs and objectives. As a result, organizations can optimize resource allocation, minimize disruptions, and drive continuous innovation, ultimately delivering greater business value and competitive advantage.

This means nothing less than the following: by leveraging the potential of data streaming, organizations position themselves to foster a culture of innovation and continuous improvement.

2.3 Observability

As IT infrastructures grow more complex, observability is becoming increasingly important. In the software context, this refers primarily to the ability to monitor or measure the internal state of a system by examining its outputs, such as logs, metrics, and traces.

- **Log aggregation** centralizes logs from various system components into a single platform, allowing for unified analysis and troubleshooting by correlating events across the system.
- **Metrics** involve the collection and monitoring of quantitative data points, such as CPU usage or request latency, providing real-time insights into the health and performance of a system.
- **Distributed tracing** tracks requests as they traverse through different services in a distributed system, offering a detailed view of the execution flow to identify performance issues and dependencies.

Observability enables developers and operators to understand, diagnose, and improve system performance and behavior in real time, which is crucial for maintaining reliable and efficient software systems.

Relation of Observability to Data Streaming

A data streaming platform has the potential to serve as a foundation for observability efforts. In an event-driven architecture, logs, metrics and traces are continuously ingested, processed, and analyzed in real time, enabling immediate insights and actions using stream processing. This real-time capability is essential for handling the high volume and high throughput data generated by modern systems.

Data streaming is complementary to batch-based log analytics tools and distributed tracing. A data streaming platform handles the real-time data pipeline to ensure that logs, metrics, and traces are processed and filtered before being ingested into an analytics platform for long-term storage and detailed analysis using indexing, searching, and visualizing the events.

Stream processing as part of the data streaming platform is best suited for time-sensitive and operational critical events because it provides fault-tolerant continuous processing in low latency, even at large scale with gigabytes per second throughput. For other use cases, the data products can be reused by simply connecting another data sink. The event-driven architecture enables reusability of data products, data pluggability with policy enforcement regarding the data contract, and ease of implementation with fast time to market.

Observability with Data Streaming:
Understand, Diagnose, and Act Continuously

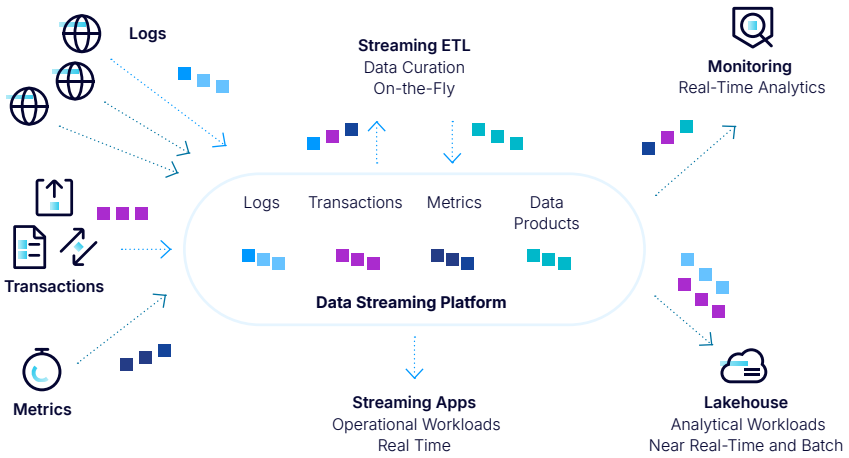

Observability with data streaming: Understand, diagnose and act continuously

Examples for Observability with Data Streaming

Consider a technology company that uses observability to monitor the performance and health of its systems: By integrating logs and metrics from various sources, such as application servers, databases, and network devices, the company can gain a comprehensive view of its infrastructure. This real-time data is then used to identify and resolve issues before they impact users, ensuring high availability and performance. A log analytics platform complements data streaming to index the log data for long-term storage, dashboard analysis, and forensics.

Another example is an e-commerce company that uses clickstream analysis to understand user behavior and optimize its website and mobile app. By leveraging data streaming, the company can continuously ingest and process clickstream data from its website, to provide real-time insights into user interactions. This information is then used to personalize the user experience, improve the shop website performance, and increase conversion rates.

In the context of cybersecurity, a financial institution might use data streaming to monitor and analyze log data for potential threats. By continuously ingesting and processing logs from firewalls, intrusion detection systems, and other security devices, the institution can detect and respond to threats in real time. This proactive approach helps to mitigate risks and protect sensitive data.

Security requires a data-centric approach where source data is collected from any environment it is created in. It is published one time, enriched into data products using any number of tools, and then made available for any security function that needs the data to be used with the tools they prefer.

Observability Through Data Streaming for Optimized Reliability and User Satisfaction

Observability through data streaming provides significant business value by enabling real-time monitoring and analysis of system events, allowing organizations to detect and respond to issues instantly, minimizing downtime and improving operational efficiency. This approach enhances visibility across complex IT environments, supports proactive maintenance, and strengthens security. Results are better service reliability and user satisfaction by enabling rapid identification of anomalies or threats.

Streaming aggregation of events also enables real-time user behavior analysis by capturing and processing clickstream data instantaneously, allowing businesses to personalize user experiences, track and trace business events across the supply chain, and respond swiftly to changing user preferences.

Data Centric Security

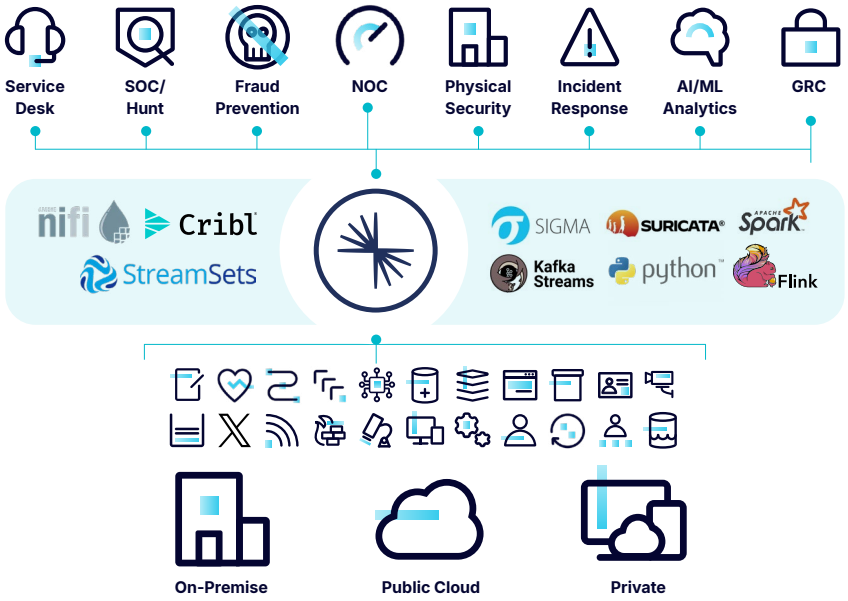

Data centric security with a data streaming platform

2.4 Hybrid and Multi-Cloud

Data replication is a technique used to replicate data across multiple data centers located in different geographic regions across on-premise, hybrid, and multi-cloud environments. This ensures that data is synchronized, allowing each data center or cloud region to have an up-to-date copy of the data, which is crucial for disaster recovery and high availability. In an active-active setup, data replication allows all data centers to be operational simultaneously, providing access to data for users regardless of their location. This setup enhances system reliability by enabling data continuity even if one data center or cloud provider experiences an outage. Additionally, cross data center replication supports load balancing to improve performance, resilience, and cost-efficiency in global applications.

Data streaming is often better for cross data center replication than traditional databases or data lakes because it enables real-time, continuous replication of data in motion with lower latency and greater flexibility. A data streaming platform efficiently handles large volumes of data and can replicate not just static data but also dynamic, time-sensitive event streams to ensure that all data centers are always in sync. This approach is particularly convenient for applications that require immediate updates and real-time processing across distributed locations, where traditional database replication might struggle with latency or scalability.

The data streaming platform facilitates this uni- or bidirectional replication and enables seamless data flow between on-premise, edge, and cloud environments by natively using the streaming technology to link clusters together for the real-time synchronization. No additional tools or infrastructure are needed.

Stream processing further enhances data replication by enabling the preprocessing of data before replication. This includes filtering, aggregating, and anonymizing data, ensuring that the data replication cost is minimized and that only relevant and compliant data is replicated to other regions, clouds, or specific applications.

Relation of Hybrid and Multi-Cloud to Data Streaming

While databases store data at rest and replicate in batch processes, data streaming provides the backbone for real-time data replication and synchronization across multiple data centers and clouds. In an event-driven architecture, data is continuously ingested, processed, and replicated across multiple data centers to ensure that information is always up to date and consistent. This real-time synchronization is crucial for maintaining data quality and enabling seamless integration between different environments.

Hybrid and multi-cloud replication across data centers with data streaming

Examples for Hybrid and Multi-Cloud with Data Streaming

Imagine an edge-to-cloud aggregation scenario in the manufacturing process across the logistics supply chain. By using data streaming, the factories and transportation vehicles can continuously ingest data from various sensors and equipment at the edge, e.g., temperature measurements, aggregate it every few minutes, and replicate only the relevant information to a cloud-based analytics platform. This enables real-time monitoring and analysis of supply chain operations for continuous optimization.

In a multi-cloud disaster recovery scenario, a company might need to replicate data between two cloud providers to ensure business continuity even if a cloud vendor has an outage. By leveraging data streaming, the company can continuously replicate data between these cloud environments to ensure that information is always available and up to date, even in the event of a failure of a cloud provider.

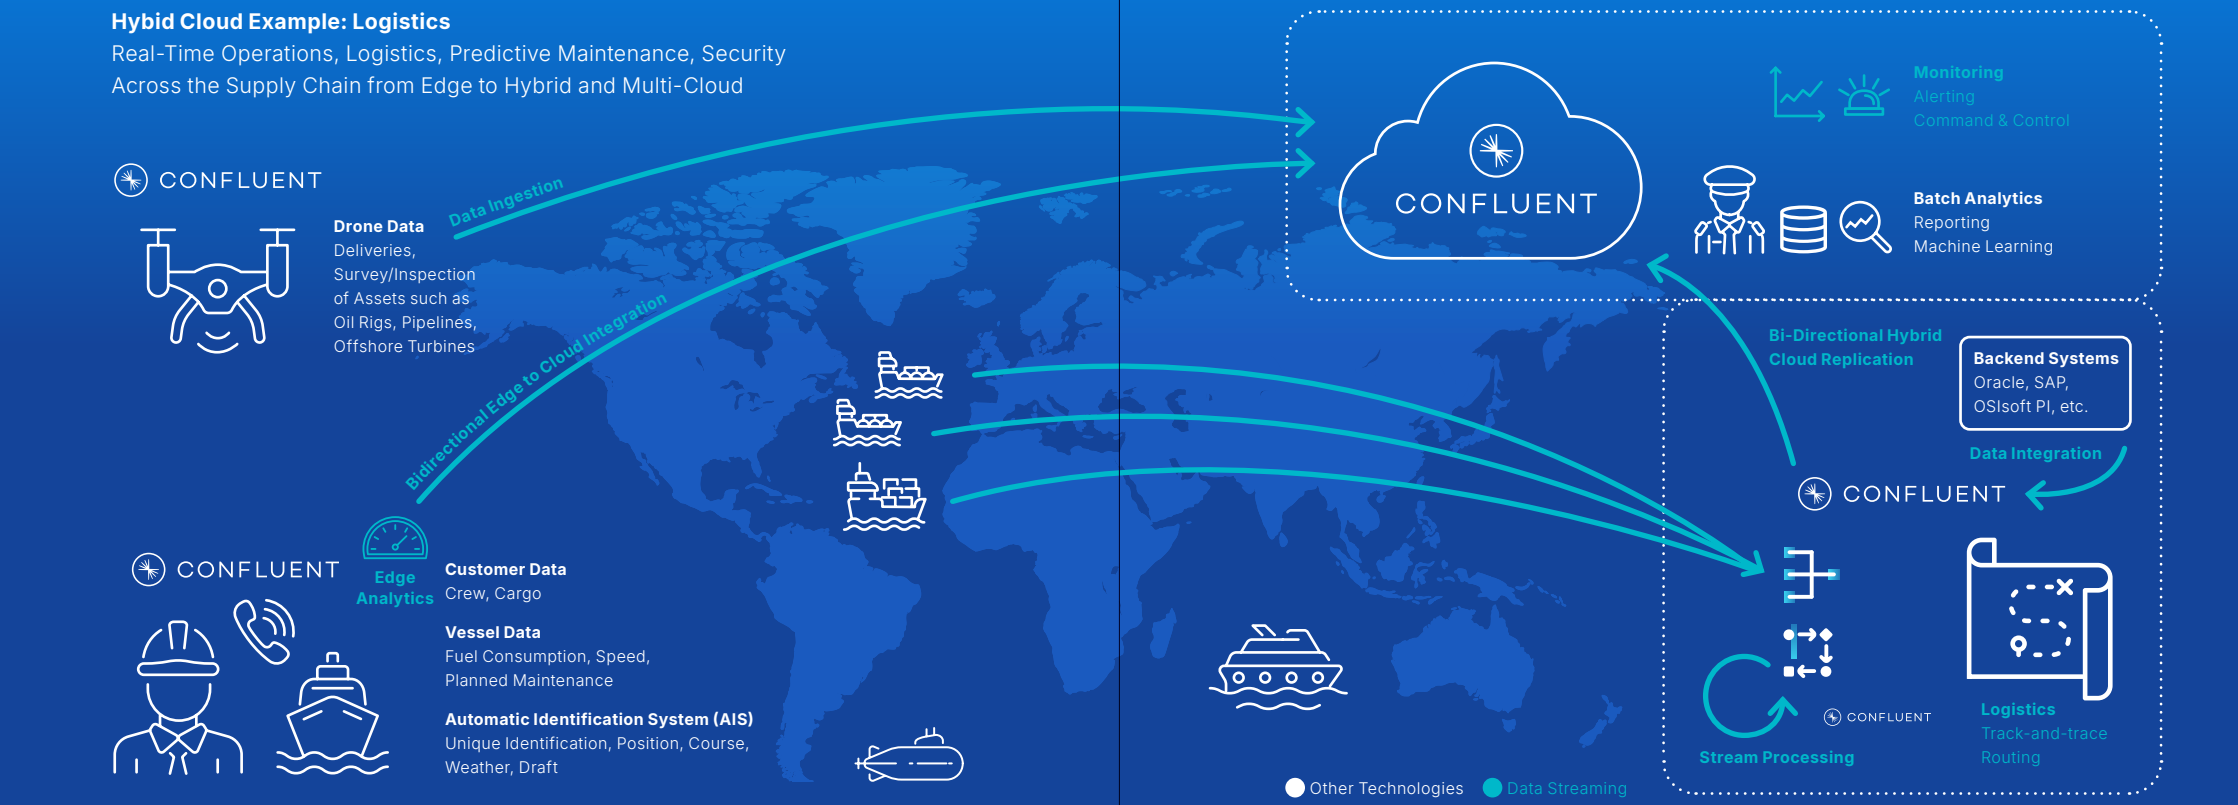

For a migration scenario, consider a company transitioning from an monolithic on-premise ERP system to a cloud-based elastic and flexible SaaS ERP. Data streaming helps the company to continuously replicate data between these systems to ensure a smooth and seamless migration process without time pressure. The final cutover (i.e., shutting down the on-premise application) happens after all applications are migrated successfully.

In a global deployment scenario, a software company might need to ensure regional data processing for worldwide availability of its application. By leveraging data streaming, the company can continuously replicate relevant data between different regions to ensure that information is always up to date and consistent, regardless of the user's location.

Hybrid and Multi-Cloud with Data Streaming for Global Availability and Scalability

The implementation of a data streaming platform for cross data center replication across hybrid and multi-cloud environments represents a strategic investment in the future of an organization's IT infrastructure. By enabling real-time data replication and processing, project teams can ensure that their organizations remain agile, responsive, and well-positioned to capitalize on emerging opportunities. This holistic approach to data replication drives operational efficiency, reduces costs, and fosters a culture of innovation and continuous improvement.

2.5 Event-Driven Microservices

A microservice architecture structures an application as a collection of small, independently deployable services, each focused on a specific business function. These services communicate over lightweight protocols and share data via data products. Microservices enable greater flexibility, scalability, and resilience compared to monolithic applications.

Building microservices with domain-driven design as a software development approach combined with an event-driven architecture is perfect for writing decoupled applications. This allows each business unit to choose its own technology stack. From a technology perspective, an event-driven microservice can be a business service, legacy monolith, external SaaS service, or an API interface.

Relation of Event-Driven Microservices and Data Streaming

While the first microservice architectures were built with point-to-point APIs, the software community realized that data streaming is ideal as the foundation of microservices because it enables real-time communication and data sharing between truly decoupled services. The event broker ensures that each microservice can independently process and respond to events as they happen. This approach supports decoupled, asynchronous interactions, which enhances the scalability and resilience of microservices architectures by allowing services to operate and evolve independently.
Additionally, data streaming ensures that microservices can maintain a consistent and up-to-date view of shared state, which is crucial for coordinated processing in distributed systems.

While the data sharing in a microservice architecture is event-based, each business unit can choose its own technology stack, whether it is a lightweight microservice, a legacy monolith, an external SaaS, or a request response API. The data streaming platform enables seamless data sharing between these diverse technologies and communication paradigms and ensures that information is always accurate and up to date.

An underestimated additional benefit of the true decoupling of applications using a data streaming platform: The event persistence layer of the data streaming platform allows travel back in time to replay and process historical events with guaranteed ordering.
This is a key differentiator from tightly coupled traditional integration platforms like message queues, ETL, ESB, iPaaS, and API management.

Event-Driven Decoupled Microservices
(Real Time, Batch and Request-Response)

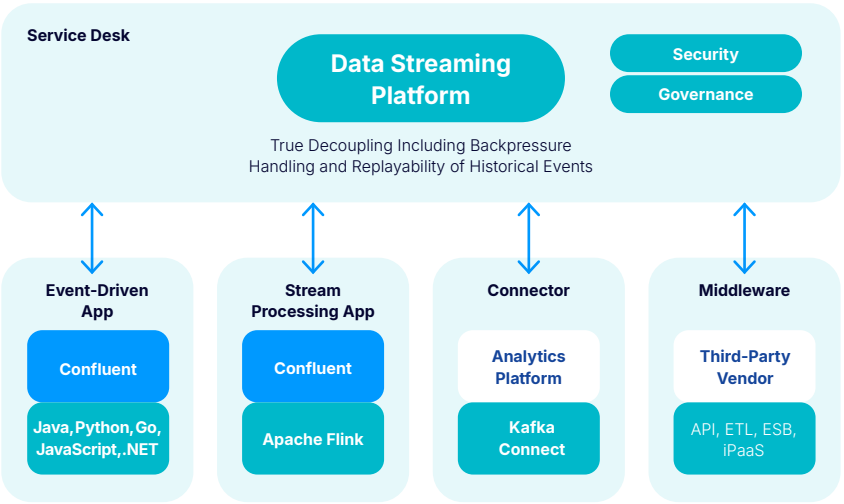

Data streaming as foundation for event-driven decoupled microservices using real-time, batch, or request-response communication

Examples for Event-Driven Microservices and Data Streaming

Consider a retail company that uses an event-driven architecture to manage its inventory and sales data. By leveraging data streaming, the company can continuously ingest data from various sources, such as point-of-sale systems, e-commerce stores, and supplier databases. This data is shared continuously between different microservices, such as inventory management, sales forecasting, and customer relationship management (CRM) to ensure that information is always up to date and consistent, no matter if the consumer is a worker in the warehouse with a mobile app using an API request response call, an automated real-time order management system, or a batch analytics platform creating reports at the end of the day.

"Organizations can remain agile, responsive, and well-positioned to capitalize on emerging opportunities by adopting data streaming for building and integrating microservices."

Another example is a financial institution that uses the event-driven architecture to monitor transactions and detect fraudulent activities. By integrating data from multiple sources, such as transaction logs, customer profiles, and external fraud databases, the institution can create a comprehensive view of each transaction. This data is then shared in real time between different microservices used for business applications, such as fraud detection, risk management, and customer service, enabling the institution to respond quickly and mitigate risks. In parallel, another microservice for regulatory reporting can replay and analyze historical events with a stream processor for root cause analysis if an incident happened.

Event-Driven Microservices Using Data Streaming for an Evolving Architecture and Technology-Independent Integration

Organizations can remain agile, responsive, and well-positioned to capitalize on emerging opportunities by adopting data streaming for building and integrating microservices. This holistic approach to event-driven applications drives operational efficiency, reduces costs, and fosters a culture of innovation and continuous improvement. Event-based microservices allow organizations to build highly scalable, flexible, and resilient systems that can respond quickly to changing business needs. Additionally, it allows teams to develop, deploy, and update services independently, accelerating innovation and reducing time to market for new features.

2.6 Data Sharing

Data sharing is the process of exchanging or providing access to data among different individuals, organizations, or systems. Within an organization, data sharing facilitates collaboration and decision-making by ensuring that all departments have access to the necessary information.

Externally, data sharing can enhance partnerships and drive innovation by allowing organizations to collaborate with external entities. The OpenAPI approach enables new business models by providing standardized interfaces for data exchange, allowing businesses to easily integrate and share data with partners and third-party applications for innovation and monetization.

Relation of Data Sharing and Data Streaming

Data streaming plays a crucial role in modern data sharing by enabling real-time data exchange between producers and consumers. With a data streaming platform, data sharing becomes seamless, whether the application is real time, batch, or request-response and whether the application uses the streaming protocol, an API, or a connector.

Consumers always have access to consistent data, regardless of the data's origin, or the time it was produced or what interface was used. Data streaming

Data Sharing (Inter-Domain, Multi-Domain, External)

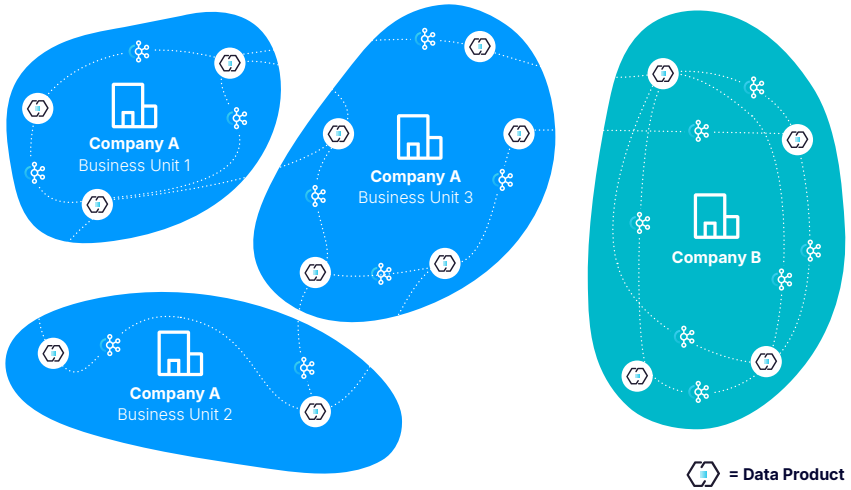

Data streaming as foundation for event-driven decoupled microservices

allows the creation of data products that can be easily shared across different applications that can use any technology to provide a unified view of data. This is particularly beneficial in event-driven architectures, where data is continuously generated and consumed by various services and applications. A data streaming helps implement efficient data sharing strategies that support both internal and external data exchange.

Data sharing can evolve to data streaming clusters across regions, multiple clouds, or hybrid environments. Unlike traditional data sharing methods, which often involve complex and time-consuming processes, data streaming allows for seamless data exchange across different environments. In a multi-cloud or hybrid setup, data streaming enables organizations to maintain data consistency and availability, regardless of the underlying infrastructure. This flexibility is essential for modern businesses that operate in diverse and dynamic geographical and legal environments.

Examples for Data Sharing

A prime example of data sharing is the financial services industry, where data sharing is critical for risk management and fraud detection. By using a data streaming platform, financial institutions can share data across different departments in real time. This enables them to quickly identify and respond to potential threats to improve the overall security and compliance.

Another example from the IoT space is a real-time data exchange between a car manufacturer and suppliers, as well as mobility services. This system allows better collaboration to publish and receive information about current orders, demand forecasts, and building new business models via partnerships. By leveraging data streaming, all stakeholders have access to the most up-to-date information, enabling timely and informed decision-making.

This real-time data flow ensures that manufacturers can adjust production schedules dynamically in response to fluctuating market demands, thereby optimizing resource allocation. Suppliers benefit from this system by receiving immediate updates on inventory requirements, reducing the risk of overproduction or stockouts. Additionally, mobility services can integrate real-time vehicle data to enhance service offerings, such as predictive maintenance and personalized customer experiences.

Data Sharing Example:
Internal and Cross-Company in Automotive

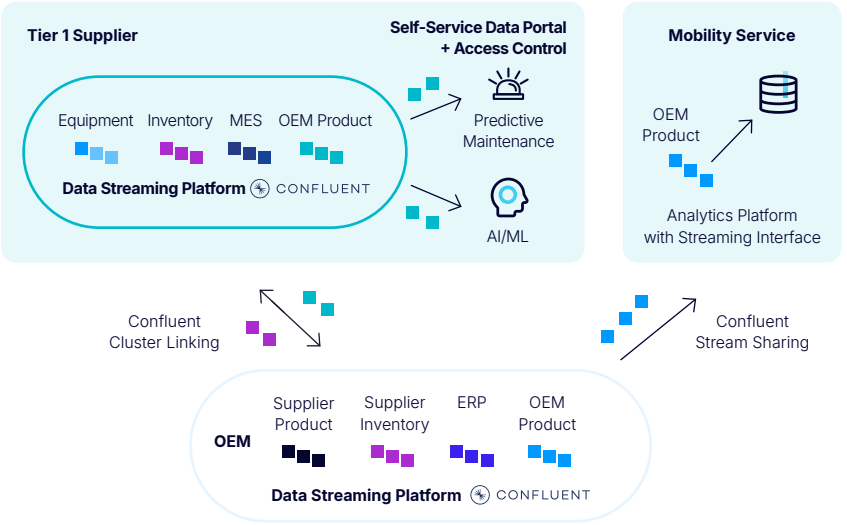

Continuous data sharing between multiple organizations

Data Sharing Through Data Streaming for Consistent Data Accessibility in Real Time

The business value of data sharing lies in its ability to enhance collaboration, drive innovation, and improve decision-making. By enabling seamless data exchange within and outside the organization, businesses can unlock new opportunities and create more efficient processes or better monetization of existing datasets. The OpenAPI approach for streaming data further enhances this value by providing standardized interfaces for data exchange to allow businesses to easily integrate with partners and third-party applications.

Data sharing allows organizations to accurately allocate costs to internal cost centers and develop external data-driven business models, driving revenue growth through optimized resource usage and new monetization strategies. This capability is integrated into the data streaming platform, enabling real-time data flows for enhanced cost management and business innovation.

2.7 Artificial Intelligence (AI)

Artificial intelligence (AI) and machine learning (ML) are transformative technologies that empower systems to learn from data and make informed decisions. Machine learning is a subset of artificial intelligence that enables systems to automatically learn from data and improve without being explicitly programmed.

Predictive AI leverages historical data to anticipate future occurrences, enabling businesses to optimize operations and enhance decision-making processes. By analyzing patterns and trends, predictive models can forecast outcomes such as customer behavior, market trends, and potential risks like equipment failure.

On the other hand, generative AI concentrates on crafting new content by assimilating existing data to generate outputs such as text, images, and designs.

Both predictive and generative AI play crucial roles in advancing automation and enhancing human capabilities across various industries. As these technologies continue to evolve, they offer unprecedented opportunities for innovation creating new business models and improving the efficiency of existing business processes.

Relation of AI and Data Streaming

Data streaming is essential to AI use cases by providing the backbone data ingestion into AI platforms, data processing to create good data quality before the ingestion, model inference in real time for critical and low-latency scenarios, monitoring and observability of the end-to-end AI/ML infrastructure. As the event-driven data streaming platform continuously ingests, processes, and analyzes data, it ensures that information is always up to date and accurate. This capability is essential for maintaining data quality and enabling seamless integration between different AI applications for model training and predictions.

Data streaming improves predictive AI by providing real-time data feeds in good data quality as input that allow models to continuously update and refine their forecasts based on the latest information. Stream processing enhances the accuracy and timeliness of model predictions and provides critical SLAs and low latency.

Data streaming supports generative AI by supplying a constant stream of data that the large language models (LLM) can use to create up-to-date, contextually relevant outputs, such as generating real-time recommendations or automated content using a new design pattern called retrieval augmented generation (RAG).

Data Streaming as Orchestration Layer for Predictive AI

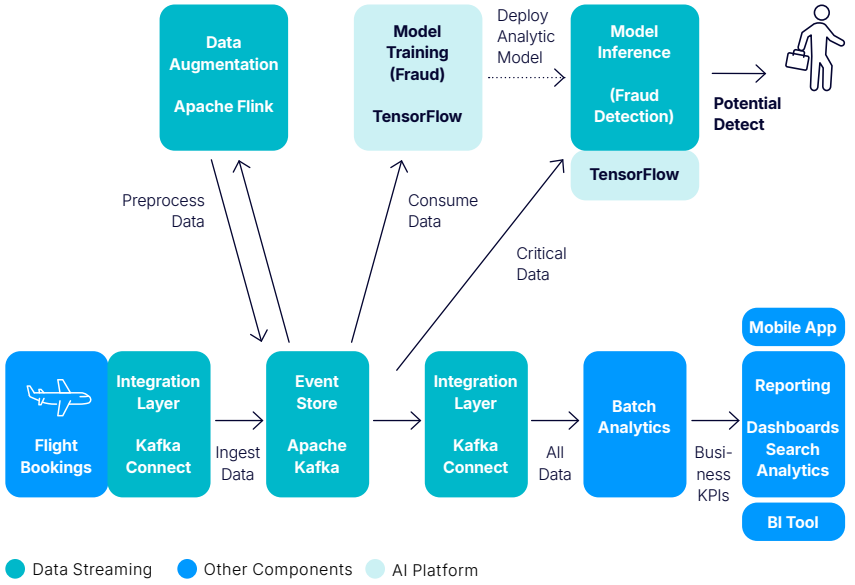

Data Streaming as orchestration layer for Predictive AI

This dynamic data flow ensures that the generative AI's outputs are highly relevant and responsive to current conditions, driving innovation and personalization in applications like content creation, interactive user experiences, or transactional decisions.

Examples for AI using Data Streaming

Consider a financial institution that uses predictive AI for fraud prevention. By leveraging data streaming, the institution can continuously ingest and process transaction data, ensuring that models are always trained on the most recent data. In parallel, risk applications apply the trained models in real time to identify suspicious patterns and trigger alerts so that the institution can respond quickly and mitigate risks. Fraud prevention is only possible if the fraud is detected before the complete business process is finished. Therefore, the model inference needs to execute during the operational payment processing, not later in the analytics platform.

Another example is a manufacturing company that uses predictive AI for predictive maintenance. By integrating data from various sensors and equipment, the company can continuously monitor the condition of its equipment and predict potential failures via real-time model inference. Alerts can then schedule maintenance activities to

Data Streaming and Retrieval-Augmented Generation (RAG) for Generative AI

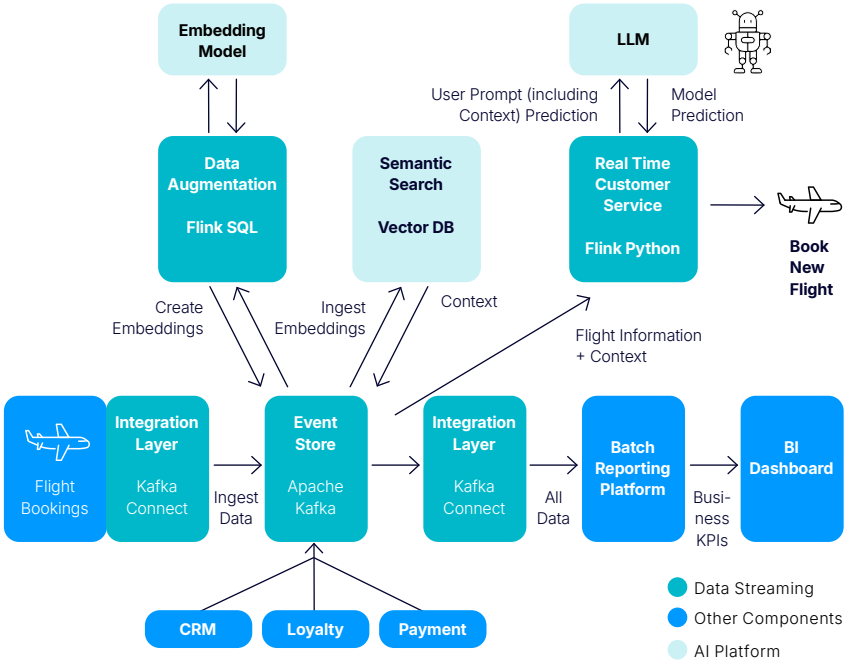

Data streaming and retrieval-augmented generation (RAG) for generative AI

reduce downtime and improve the overall equipment effectiveness (OEE).

In the context of generative AI, consider an airline that uses AI to provide customer service for flight rebookings. By integrating data from various sources, such as booking engines, flight schedules, and customer profiles, the airline can provide real-time responses to customer inquiries. The chatbot is used to generate personalized recommendations and rebooking options, improving customer satisfaction and loyalty. The GenAI application can even trigger transactional actions like a rebooking or payment process through the data streaming pipeline.

Data Streaming for More Powerful and Accurate Results in AI Use Cases

The usage of a data streaming platform tailored for AI use cases is not just a technological upgrade; it is a crucial strategy for organizations aiming to leverage real-time insights. Accurate, timely decisions are essential in fast-paced markets and can be enabled by streamlining data ingestion, processing, and analytics. The adoption of data streaming for predictive and generative AI use cases offers significant business value by enhancing the data quality, improving the model training and model inference process, and enabling real-time decision-making. Organizations can seamlessly integrate and process any operational or analytical data for AI applications to ensure that information is always accurate and up to date.

Furthermore, this approach enables the continuous training and retraining of AI models with fresh data to foster a dynamic environment where businesses can adapt quickly to changing consumer behaviors and emerging trends. Ultimately, integrating a data streaming platform allows organizations to harness the full potential of AI and leads to innovative solutions that are optimized operational processes, and significant cost savings.

Predictive AI Example: Fraud Detection

Stateless and stateful stream processing for real-time data correlation with Apache Flink

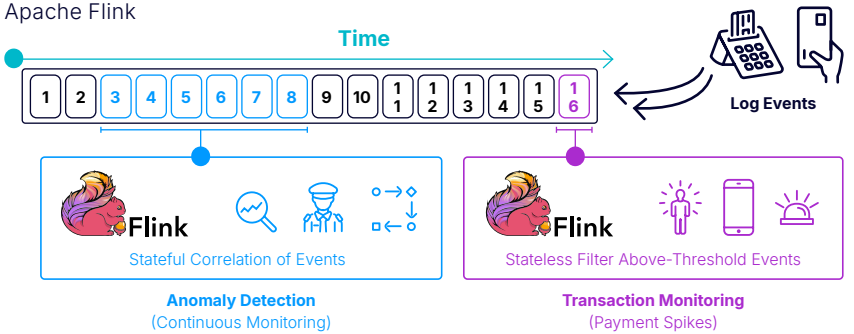

Automated and context-specific customer service in real time by using GenAI

Generative AI Example: Context-Specific Customer Service in Real-Time with Stream Processing

Stateless and stateful stream processing for real-time data correlation with Apache Flink (Flink SQL)

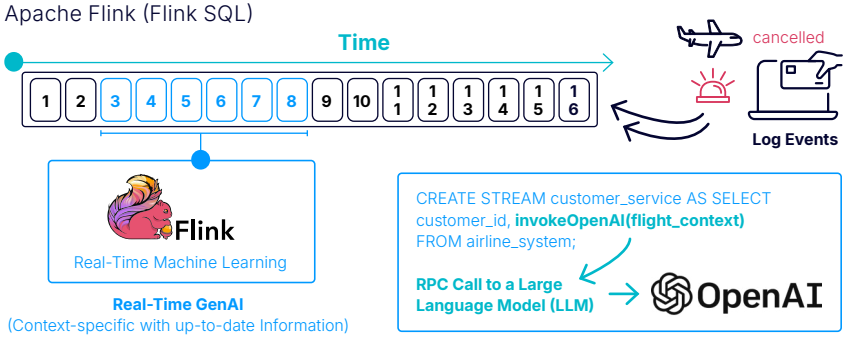

Anomaly detection for fraud prevention with continous stream processing

2.8 Internet of Things (IoT)

The Internet of Things (IoT) refers to the network of interconnected devices that communicate and share data over the internet, enabling automation and smarter decision-making across various domains.

Industrial IoT (IIoT) connects sensors, equipment, and systems in sectors like manufacturing or energy and utilities to optimize operations through real-time data and analytics. The integration of operational technology (OT) with information technology (IT) in IIoT allows for seamless data flow and advanced process control, significantly improvingthe overall equipment effectiveness (OEE) by reducing downtime, enhancing productivity, and enabling predictive maintenance.

Consumer IoT (CIoT) includes everyday devices such as smart home appliances, wearables, and connected cars to enhance convenience, personalization, and efficiency for users. These devices seamlessly integrate into daily life, offering real-time data and insights that empower users to make informed decisions and automate routine tasks. As a result, Consumer IoT not only improves the quality of life but also fosters a more connected and intelligent living environment.

Example: Industrial IoT

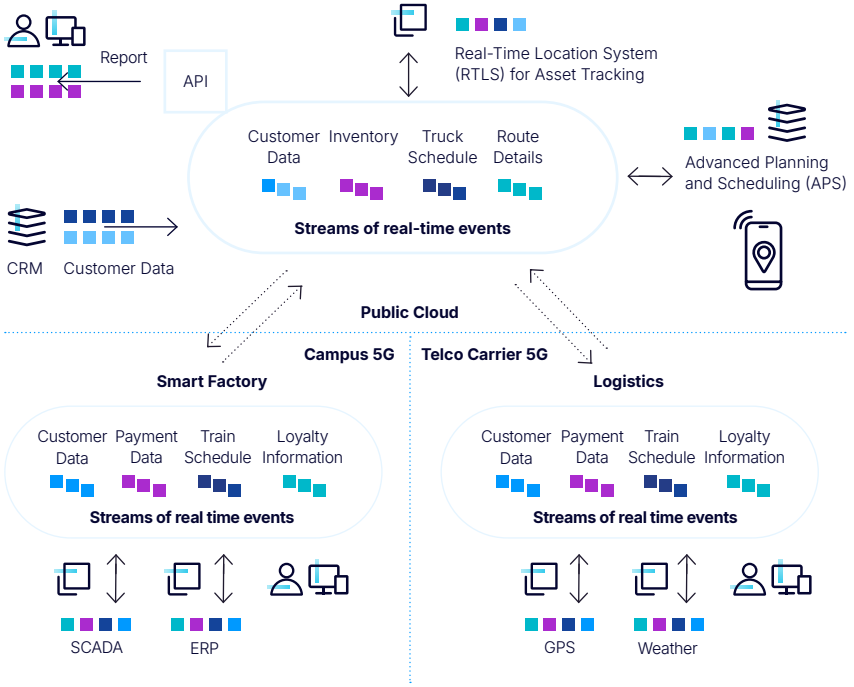

Improvement of overall equipment effectiveness (OEE) with data streaming across edge and hybrid-cloud environments

Data Streaming at the Edge in the Smart Retail Store

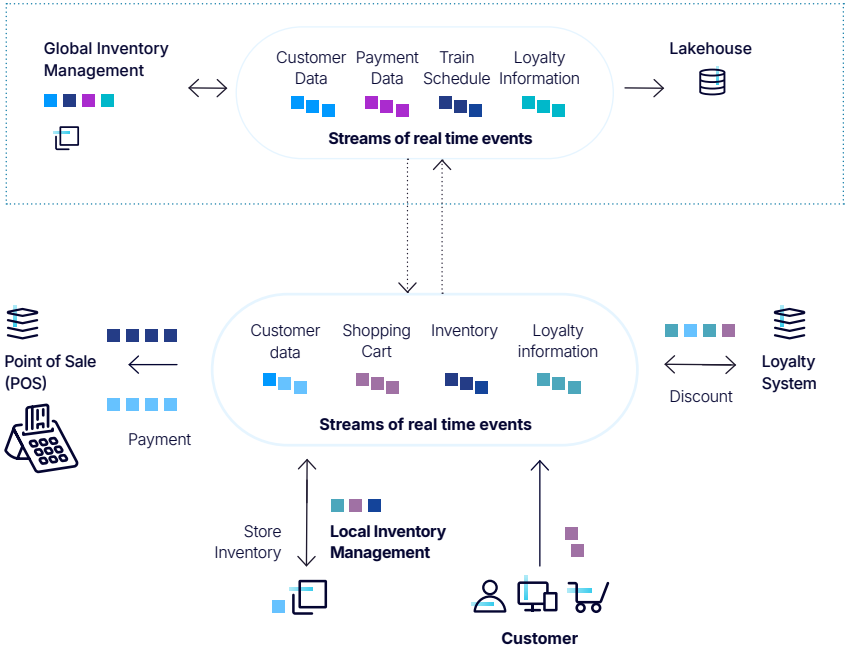

Integration and synchronization between smart retail stores and global inventory mangement in the cloud

Relation of IoT and Data Streaming

IoT naturally generates data continuously to be ingested, processed, and analyzed in real time. Data streaming is essential for managing the vast amounts of data generated by IoT. Stream processing ensures that information is continuously aggregated and actionable. While data streaming technologies such as Kafka and Flink are not hard real time (meaning deterministic communication with zero spikes and zero latency) and therefore not suitable to build safety-critical embedded systems like a robotic operating system or Advanced Driver Assistance Systems (ADAS) in a car, they are highly applicable when processing non-safety- critical data. Data streaming fits into the enterprise architecture for IoT use cases by providing a scalable and reliable platform for continuous data processing.

A hybrid edge-cloud architecture offers the flexibility to deploy applications and services where they are most needed. This flexibility is enabled by providing a consistent data pipeline that can operate across both edge and cloud environments. This allows organizations to optimize their infrastructure based on specific use cases and requirements. Many organizations choose a cloud-first approach for new projects where latency, security, and cost are the main factors to decide if a service can be deployed in the cloud or needs to be at the edge.

Examples for Data Streaming in the IoT World

In the Industrial IoT space, a manufacturing company can use a data streaming platform at the factory shop floor. This includes logistics, communication with PLCs and robotics, monitoring of equipment, and transactional decisioning using MES/ERP and related systems. By leveraging data streaming, the company can create a bridge between the OT world (usually at the edge) and the traditional IT world (usually in a private data center or the public cloud), enabling real-time monitoring and analysis of factory operations.

In a smart retail store, data streaming at the edge connects inventory management, location-based services, and point-of-sale (POS) systems to provide real-time updates and personalized customer experiences. Customer actions are quickly captured and sent to a hybrid cloud, where a loyalty platform and CRM system use the data to offer customized promotions and rewards. This setup improves customer satisfaction, keeps inventory levels just right, and boosts sales, helping the store grow its customer base.

Data Streaming Platform as Foundation for an Open and Scalable Data Historian

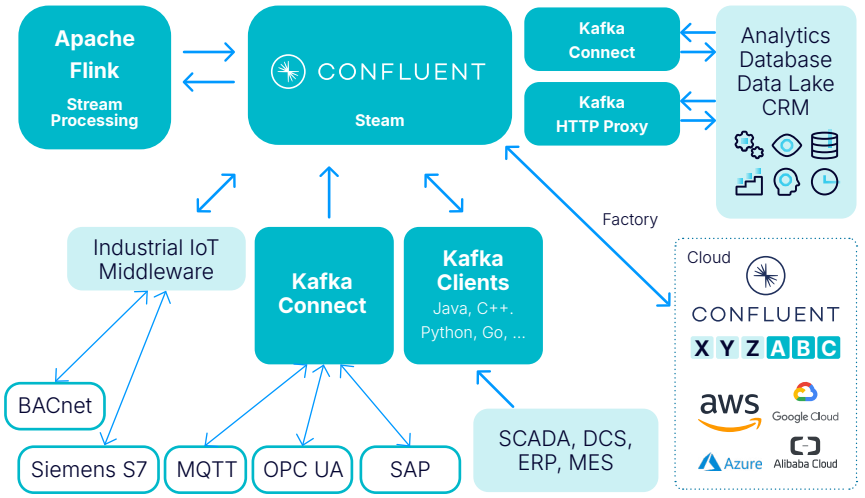

IT/OT integration between the smart factory at the edge and analytics in the cloud

IoT with Data Streaming to Optimize the OEE and Provide Innovative Consumer Experiences

Data streaming in IoT projects provides significant business value by enabling continuous data processing and decision-making. The (near) real-time capability allows organizations to monitor and respond to events as they happen, improving operational efficiency and reducing downtime.

By leveraging data streaming, organizations can seamlessly integrate and process high-volume IoT data from logs, sensors, robotics, equipment, and other interfaces, and combine it with low-volume transactional data from operational systems like MES, ERP, or CRM.

In industrial settings, data streaming enhances predictive maintenance by continuously analyzing sensor data to prevent equipment failures and optimize maintenance schedules and improve the OEE. For Consumer IoT, it supports dynamic personalization and rapid response to user behaviors, enhancing customer experiences. Additionally, data streaming enables scalable and flexible analytics, allowing businesses to process large volumes of data from numerous devices without delays, which is critical for gaining insights and driving informed decisions in fast-paced environments. Overall, it accelerates innovation, reduces costs, and enhances competitiveness by ensuring that data is actionable as soon as it's generated.

2.9 Data Governance

Data governance refers to the comprehensive management of data's availability, usability, integrity, and security within an organization. It involves establishing processes, roles, policies, standards, and metrics to ensure that data is properly managed throughout its lifecycle. The primary goal of data governance is to ensure that data is accurate, consistent, secure, and compliant with regulatory requirements and organizational policies.

- **Compliance** in data governance ensures that data management practices adhere to relevant laws, regulations, and standards. This includes maintaining data privacy, security, and integrity to avoid legal and financial repercussions.
- **Data sovereignty** refers to the concept that data is subject to the laws and governance structures within the nation it is collected. It emphasizes the need for organizations to understand and comply with the legal requirements of the jurisdictions in which they operate.
- **Data privacy and security** focus on protecting sensitive information from unauthorized access and ensuring that personal data is handled in compliance with privacy laws. This involves implementing robust security measures and privacy policies to safeguard data.
- **Data quality** involves ensuring that data is accurate, complete, reliable, and timely. High-quality data is essential for effective decision-making and operational efficiency, as it directly impacts the outcomes of data-driven processes.

Relation of Data Governance and Data Streaming

A data streaming platform enables organizations to process and analyze data continuously as it is generated, allowing for immediate actionability and decision-making. This real-time processing capabilities on top of data contracts ensures that data governance policies and controls can be applied instantaneously to address issues such as data quality, security breaches, or compliance violations as they occur.

In an event-driven architecture, a data catalog and data lineage built on top of data streaming facilitates the continuous flow of data across various systems. This ensures that data governance mechanisms are consistently enforced. Data streaming provides the scalability and resilience needed to handle high volumes of data to assure that governance processes can keep pace with the speed and scale of streaming data environments. Built on data products with data contracts for structure, metadata, policies, integrity constraints and schema evolution, this integration allows for the automation of governance processes, enabling timely detection, notification, and remediation of governance issues.

Data Governance for Compliance, Data Souvereignty, and Quality

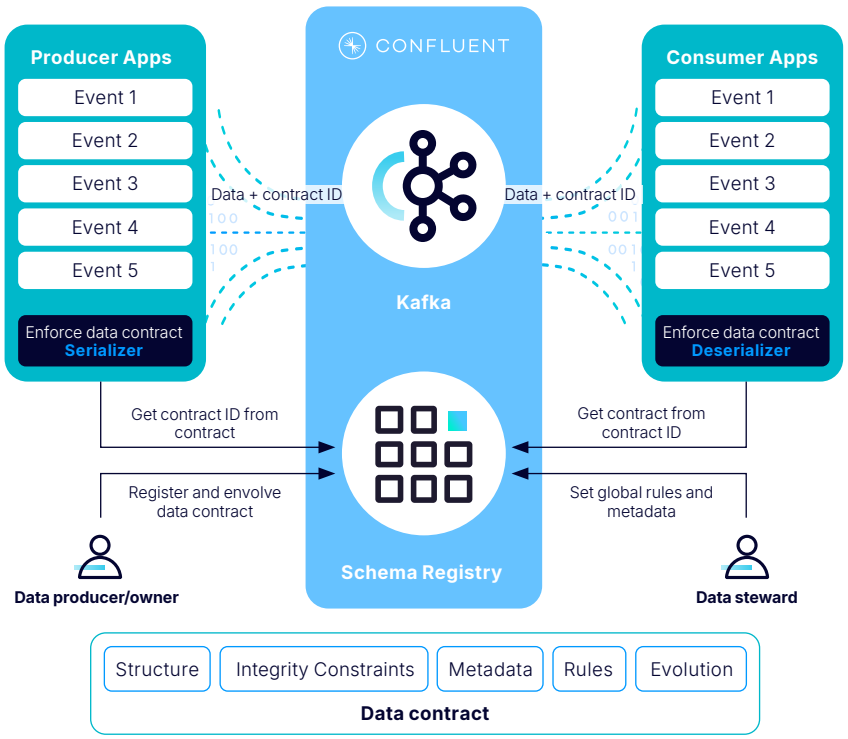

Continuous data governance for data in motion to ensure compliance, data sovereignty, and data quality

Examples for Data Governance in Streaming Data Architectures

Consider a financial services company that uses data streaming to monitor transactions in real time. By leveraging a data streaming platform, the company can implement data governance policies that detect fraudulent activities as they occur to ensure compliance with financial regulations and to protect customer data. Another example is a healthcare organization that uses data streaming to manage patient data continuously. With real-time data processing on top of data contracts, the organization can guarantee data quality and privacy by applying governance controls that validate data accuracy and enforce privacy policies, thus maintaining compliance with healthcare regulations like HIPAA (Health Insurance Portability and Accountability Act).

Stream Lineage

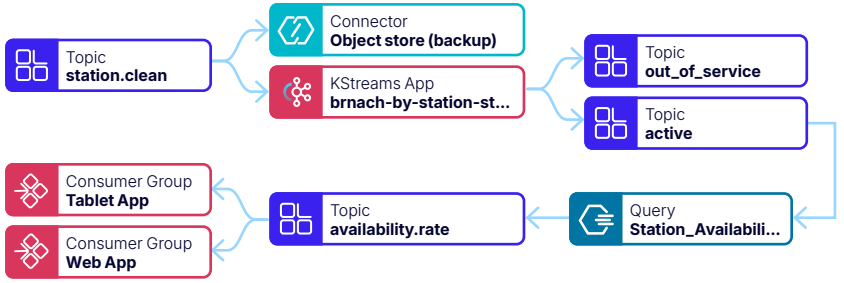

Data Catalog

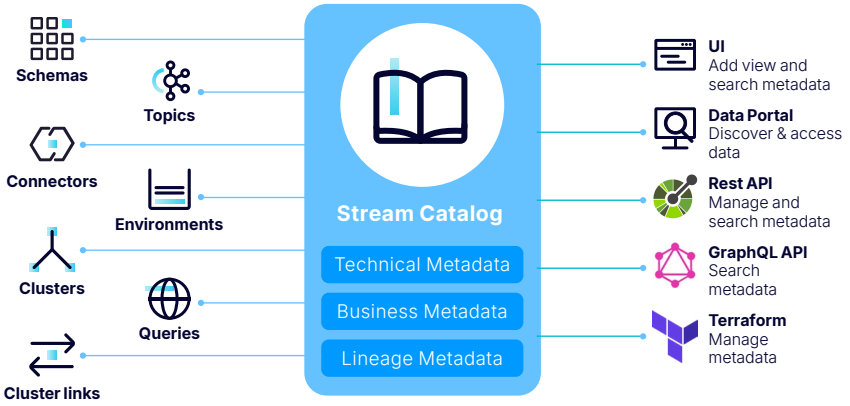

Data Portal

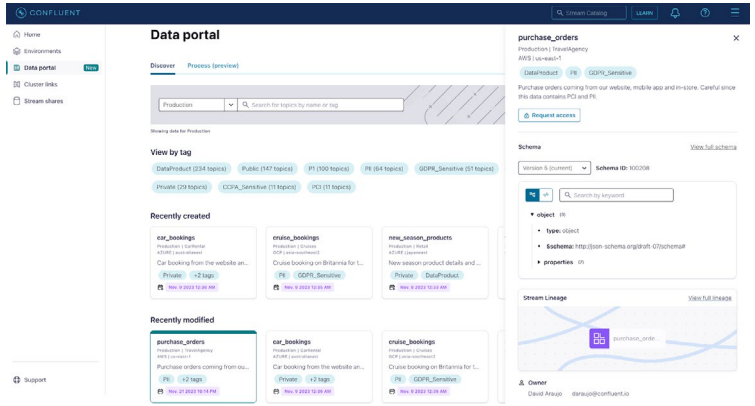

Data Governance on Top of Data Streaming for Continuous Compliance and Data Sovereignty

The adoption of data streaming offers significant opportunities for improving data governance across the entire enterprise architecture. By enabling continuous data processing and analysis, organizations can enhance their ability to manage data effectively, ensuring compliance, data sovereignty, privacy, and quality. This not only mitigates risks associated with data management but also enhances operational efficiency and decision-making capabilities. Furthermore, data streaming facilitates real-time visibility into data flows, allowing organizations to quickly identify and address potential issues before they escalate. This proactive approach to data management supports a culture of transparency and accountability, fostering trust among stakeholders. Additionally, the integration of data streaming with enterprise-wide governance frameworks ensures that data policies are consistently enforced across all platforms, further strengthening the organization's data governance posture.

2.10 ESG (Environmental, Social, and Governance)

Environmental, Social, and Governance (ESG) criteria are a set of standards for a company's operations that socially conscious investors use to screen potential investments. ESG is important across all industries, as it addresses critical issues like climate change, social responsibility, and corporate governance. Data-driven software plays a significant role in measuring, monitoring, and improving ESG performance, including regulatory requirements that are being implemented now, for instance the Corporate Sustainability Reporting Directive (CSRD).

Relation of ESG and Data Streaming

Data streaming plays a crucial role in advancing ESG initiatives b y providing real-time insights and enabling proactive management of sustainability efforts.

For environmental monitoring, continuous data streaming allows organizations to track key metrics such as air and water quality, energy consumption, and emissions in real time. This immediacy helps identify and mitigate environmental risks while optimizing resource usage and reducing carbon footprints.

In terms of supply chain transparency, data streaming ensures that companies can monitor their supply chains continuously and with good data quality, verifying that suppliers adhere to ethical and sustainability standards. This transparency improves the accuracy of ESG metrics and reporting.

On the social impact front, streaming data enables real-time monitoring of working conditions, employee well-being, and community effects, allowing companies to promptly address any social issues and ensure compliance with labor and human rights standards.

Additionally, data streaming supports governance by providing up-to-date information on regulatory compliance to help organizations adapt quickly to changing regulations and align their operations with ESG goals.

Finally, data streaming enhances the accuracy and timeliness of ESG reporting and offers stakeholders transparent and current information on a company's ESG performance, which strengthens investor confidence, regulation compliance, and corporate reputation.

Event-Driven Architecture with Data Streaming to Improve ESG

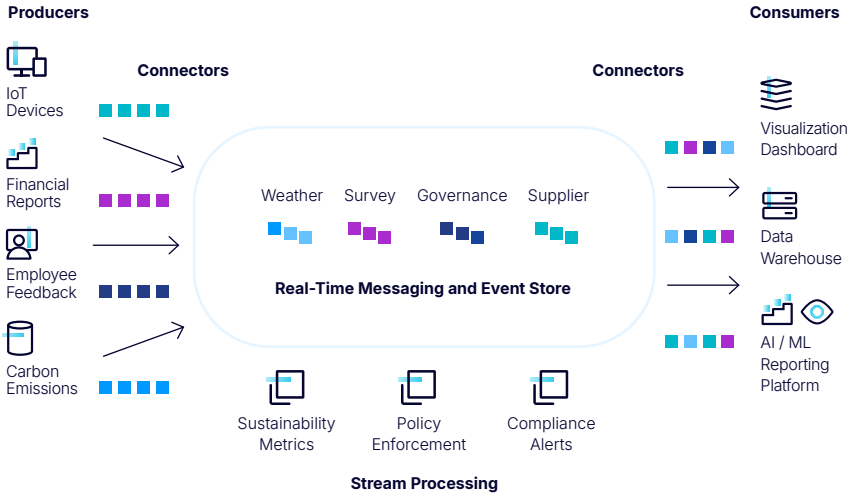

Event-Driven architecture with data streaming to improve ESG criteria

Examples for ESG Use Cases with Data Streaming

Consider a manufacturing company that uses data streaming to monitor and analyze emissions data from its industrial plants. By leveraging data streaming, the company can continuously ingest and process emissions data in real time to enable it to identify and address potential issues before they become significant problems. This real-time data processing capability is crucial for reducing the company's carbon footprint and improving its environmental performance.

In the financial services industry, data streaming can significantly enhance ESG efforts by improving transparency and risk management. For example, a bank might use data streaming to monitor and analyze real-time transactions and activities across its investment portfolio. By continuously streaming data on the environmental impact, social practices, and governance structures of the companies in which it invests, the bank can more effectively assess ESG risks and opportunities.

If a company in the portfolio suddenly faces a social or environmental controversy, the bank's data streaming system can instantly flag this issue. This allows the bank to respond quickly, such as by adjusting its investment strategy or engaging with the company to address the problem. Additionally, real-time data can be used to ensure that the bank's lending practices are aligned with ESG criteria, such as avoiding financing for projects that harm the environment or violate human rights.

Improved ESG Compliance and Alignment Through Data Streaming

Using data streaming for ESG initiatives delivers substantial business value by enabling real-time monitoring, risk management, and improved decision-making, which are critical for maintaining compliance and enhancing corporate reputation. Continuous data streaming allows companies to track and respond to ESG-related events as they happen, reducing the likelihood of environmental, social, or governance-related incidents that could lead to financial penalties, reputational damage, or regulatory scrutiny. This proactive approach to managing ESG risks can lead to cost savings by preventing issues before they escalate and ensuring compliance with evolving regulations.

Furthermore, data streaming provides accurate and up-to-date ESG metrics to improve the quality of ESG reporting. This transparency not only meets the demands of regulators but also enhances investor confidence, attracting ESG-focused investments and potentially lowering the cost of capital. By integrating real-time ESG data into decision-making processes, companies can identify new opportunities for sustainability. For instance, optimizing resource use or identifying sustainable investments, which can drive long-term growth and competitive advantage. Overall, data streaming enhances a company's ability to manage ESG factors effectively, aligning business practices with societal as well as regulatory expectations and creating value for both shareholders and stakeholders.

Industry Success Stories for Data Streaming

03

The previous chapter showcased the wide array of technical use cases where data streaming can add significant value, regardless of industry. This book aims to demonstrate that these applications are already being implemented across various sectors. While not all technical use cases were covered in Chapter 2, and not every scenario in every industry can be addressed, the focus remains on some areas with strong data streaming adoption, illustrated through real-world examples to provide maximum value to readers.

This section offers a comprehensive overview of trends across all major industries, along with concrete examples of data streaming usage by customers and highlighting the business value they gain. The aim is for both experienced users and "streaming newcomers" to explore practical applications of data streaming through specific scenarios and industry examples. All industry trends are based on current research and discussions with existing Confluent customers. The scenarios are directly linked to these trends and are evident in real customer examples.

The following subchapters include sections that present scenarios and real-life examples which address common challenges faced by institutions from specific industries (financial services, manufacturing, retail etc.) in relation to digitization and the transformation of their business models.

This chapter is just a snapshot on the many success stories and business results that are coming from data streaming. We aim to extend this book with more examples in every release.

For those interested in being featured in the next edition of "The Ultimate Data Streaming Guide," contact the authors via use-case-book@confluent.io

3.1 Financial Services

The financial services industry is undergoing a profound transformation driven by digitalization and innovation. The industry is experiencing a constant need to balance customers' high expectations for ease of use, service orientation, and availability with the traditionally high standards of security and system reliability. A rapid adoption of digital technologies is reshaping how financial institutions operate, interact with customers, and deliver services.

One of the most significant trends is the transition to cloud-based solutions, which offers unparalleled scalability, flexibility, and cost-efficiency. Hybrid cloud integration, in particular, is gaining traction as it allows financial institutions to leverage the best of both public and private clouds. This approach ensures data security and compliance while enabling seamless access to advanced analytics and real-time processing capabilities.

Additionally, the rise of fintech companies and digital-first banks is pushing traditional financial institutions to innovate and modernize their legacy systems. Real-time data streaming and processing have become essential for delivering personalized customer experiences, detecting fraud, and optimizing operations.

The integration of artificial intelligence (AI) and machine learning (ML) into financial services is further enhancing decision-making and risk management. As regulatory requirements continue to evolve, financial institutions must also ensure that their digital transformation initiatives comply with stringent standards.

Overall, the convergence of cloud computing, data streaming, and AI/ML is driving a new era of innovation and efficiency in the financial services industry.

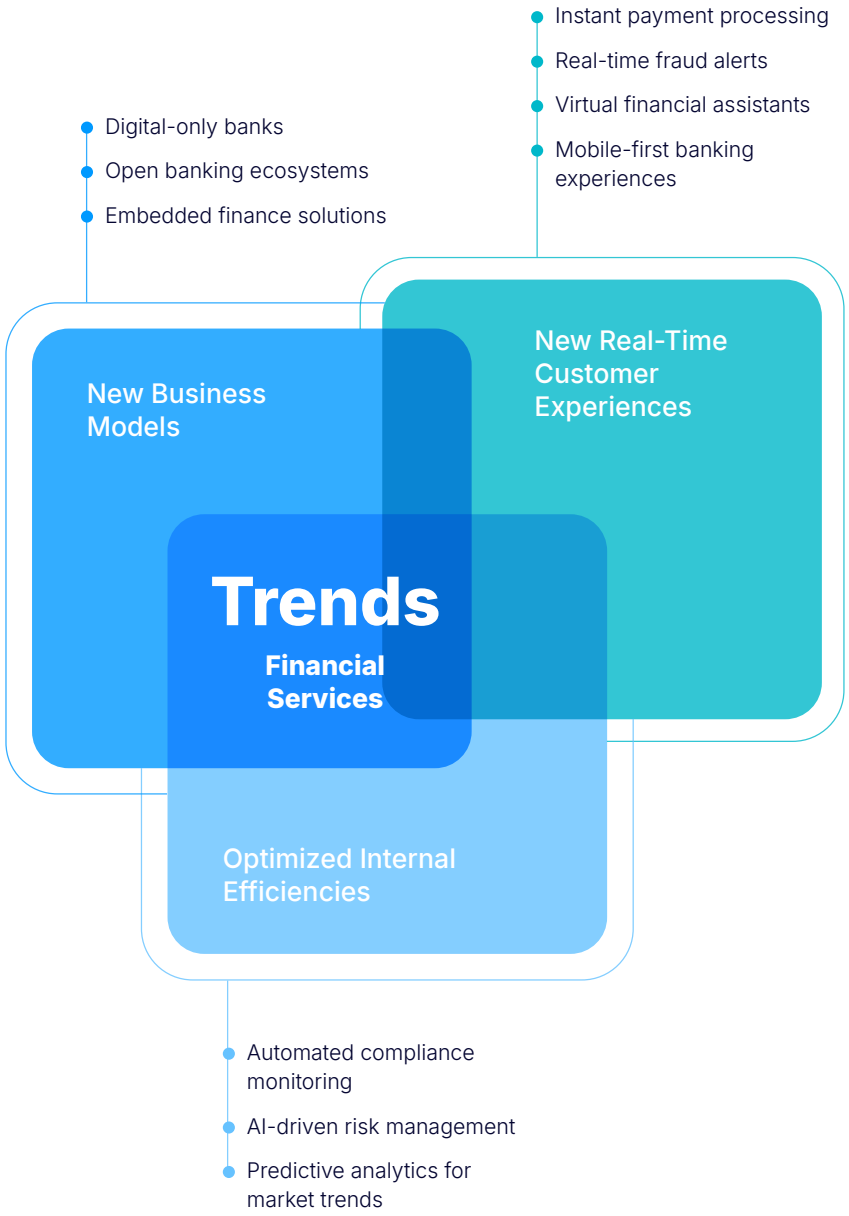

Scenario 1:
Legacy Offloading

Legacy systems, particularly mainframes, have been the backbone of financial institutions for decades, handling critical transactional workloads and storing vast amounts of data. However, the high operational costs and inflexibility of these legacy systems have become a significant burden. Mainframe offloading, using data streaming platforms, offers a compelling solution to this challenge. By migrating data and workloads to a more flexible and scalable cloud-native environment, financial institutions can drastically lower their operational expenses and enhance the agility of their IT operations. Data streaming enables data consistency between legacy batch infrastructure and modern real-time cloud streaming applications and APIs, ensuring seamless integration and real-time data availability.

Business Value

- **Cost reduction:** The transition allows banks to reduce the number of reads on the mainframe, significantly cutting down on fixed infrastructure costs (OPEX).
- **Compliance and innovation:** Financial institutions maintain compliance with regulatory requirements and business logic while enabling the creation of new applications using the same event-based architecture.
- **Operational efficiency:** The new architecture enhances the agility of IT operations, allowing for faster deployment of new services and ensuring data consistency across mainframe, databases and customer-facing applications.

"Banking systems are typically a mix of cloud-based applications and legacy databases. In scenarios where our developers don't have access to the underlying codebase for systems we want to connect to, we really like using pre-built connectors from Confluent. We can directly connect to databases, capture changes that are happening, create events from those changes and keep these systems in sync with our downstream systems."

Badr Alduaiji, *Head of Integration and Architectur, Arab National Bank*

Scenario 2:
Core Banking

Core banking systems are essential for processing transactions, facilitating payments, and managing customer accounts. Modernizing these systems to leverage data streaming platforms can significantly enhance their capabilities, providing real-time processing and greater scalability. By integrating data streaming technologies, financial institutions can build resilient and compliant infrastructures that handle high volumes of transactions with zero data loss. This modernization also enables faster time to market for new use cases and services, ensuring that banks can stay ahead in a competitive landscape.

Business Value

- **Resilient and compliant infrastructure:** The platform provides a resilient and compliant infrastructure that can handle high volumes of transactions with zero data loss.
- **Elastic scalability:** The elastic scalability of the platform ensures that banks can efficiently manage peak loads without compromising performance.
- **Faster time to market:** The ability to rapidly deploy new use cases and services accelerates time to market, enabling banks to stay ahead in a competitive landscape.

"Not only does Confluent provide lower latency and higher throughput than our legacy system, it is also more resilient."

Sunarto Rahardja, *SVP and Head of Software Development at Singapore Exchange (SGX)*

Scenario 3:
Real-Time Payments

In financial services, instant payments and high-speed trading are crucial to meet customer demand for immediate transaction processing and to keep pace with modern markets. Financial institutions must ensure that transactions, asset transfers and trades are executed without delay, especially during peak trading periods.

Traditional batch processing systems are insufficient for today's needs, leading to bottlenecks, latency problems, and operational risks.

Data streaming enables real-time processing and analyze massive transaction volumes. It also allows continuous monitoring, instant transaction validation, anti-fraud checks, and risk analysis during the transfer process.

By ingesting data across from trading systems, customer accounts, and compliance systems, data streaming enables near-instantaneous clearing and settlement. This digitizes Electronic funds transfer (EFT), covering, payments such as direct deposits, wire transfers, ACH payments, ATM transactions, and point-of-sale (POS) transactions. High-frequency trading operations with millisecond latency allow trading platforms to update risk profiles in real time and adjust strategies dynamically based on market data.

Business Value

- **Reduced latency:** Execute trades and process EFTs in real time, meeting the demands of modern high-frequency trading and immediate fund transfers.
- **Improved system resilience:** Achieve high availability and seamless failover, ensuring operations continue uninterrupted, even during high-volume periods.
- **Scalability:** Scale dynamically to accommodate peak trading times or surges in EFT activity, maintaining performance without compromising data integrity.
- **Regulatory compliance and risk mitigation:** Meeting T+0 reporting requirements ensures immediate compliance with financial regulations, reducing the risk of penalties and enhancing trust with regulators and stakeholders.
- **Faster time to market:** Deploy new features and compliance updates quickly with microservices, allowing for rapid response to regulatory changes and market demands.

Scenario 4:
Fraud Detection

Fraud detection is a critical priority for financial institutions, demanding robust, real-time analysis to detect, prevent, and respond to suspicious activities swiftly. Data streaming platforms empower institutions to implement real-time risk and anti-money laundering (AML) compliance models that continuously monitor and analyze transaction data for unusual patterns, behaviors, and anomalies that could indicate fraud or money laundering.

With advanced data streaming technologies, institutions can aggregate data from multiple sources, including transactions, customer profiles, and regulatory feeds, to generate a holistic view of potential risks. This real-time, cross-functional analysis allows for instant alerts, rapid investigation, and compliance with AML regulations, significantly reducing the window of opportunity for fraudulent actions and safeguarding both financial assets and institutional integrity.

Business Value

- **Reduced risk:** Real-time fraud and AML detection minimizes financial losses, protecting both customer and institutional assets.
- **Regulatory compliance:** Meeting AML standards in real time reduces penalties and ensures compliance with stringent regulatory requirements.
- **Enhanced customer trust:** Proactive fraud prevention strengthens customer confidence, loyalty, and brand reputation.
- **Operational efficiency:** Real-time analysis and rapid response streamline operations, improving accuracy in decision-making and resource allocation.
- **Scalable risk management:** Scalable data streaming enables efficient monitoring as transaction volumes grow, ensuring resilience and adaptability in fraud detection processes.

Scenario 5:
Regulatory Reporting

Regulatory reporting is a critical function for financial institutions, particularly those involved in the derivatives market. It involves the use of trade data repositories and reporting service providers to ensure compliance with regulatory requirements. The accuracy and timeliness of these reports are essential for maintaining transparency and operational efficiency.

Financial services regulations increasingly demand near-instantaneous, T+0 (trade date) reporting for payments and trades, moving from the previous T+1 (next-day) standard to ensure real-time transparency and compliance.

Data streaming platforms play a vital role in enhancing the regulatory reporting process. By enabling real-time data validation and processing, these platforms ensure that the data submitted to regulators is accurate and up to date. This capability not only improves compliance but also reduces the operational burden associated with manual data processing and reporting.

Business Value

- **Transparency and real-time data validation:** Real-time processing ensures that all data is accurate and validated before submission, enhancing transparency.
- **Enhanced compliance with regulatory requirements:** Continuous data validation and processing ensure that all regulatory obligations are met promptly.
- **Operational efficiency:** Automating the data processing and reporting functions reduces the need for manual intervention, saving time and resources.
- **Competitive advantage:** By leveraging real-time data streaming, financial institutions can set themselves apart from competitors that rely on traditional, slower methods.

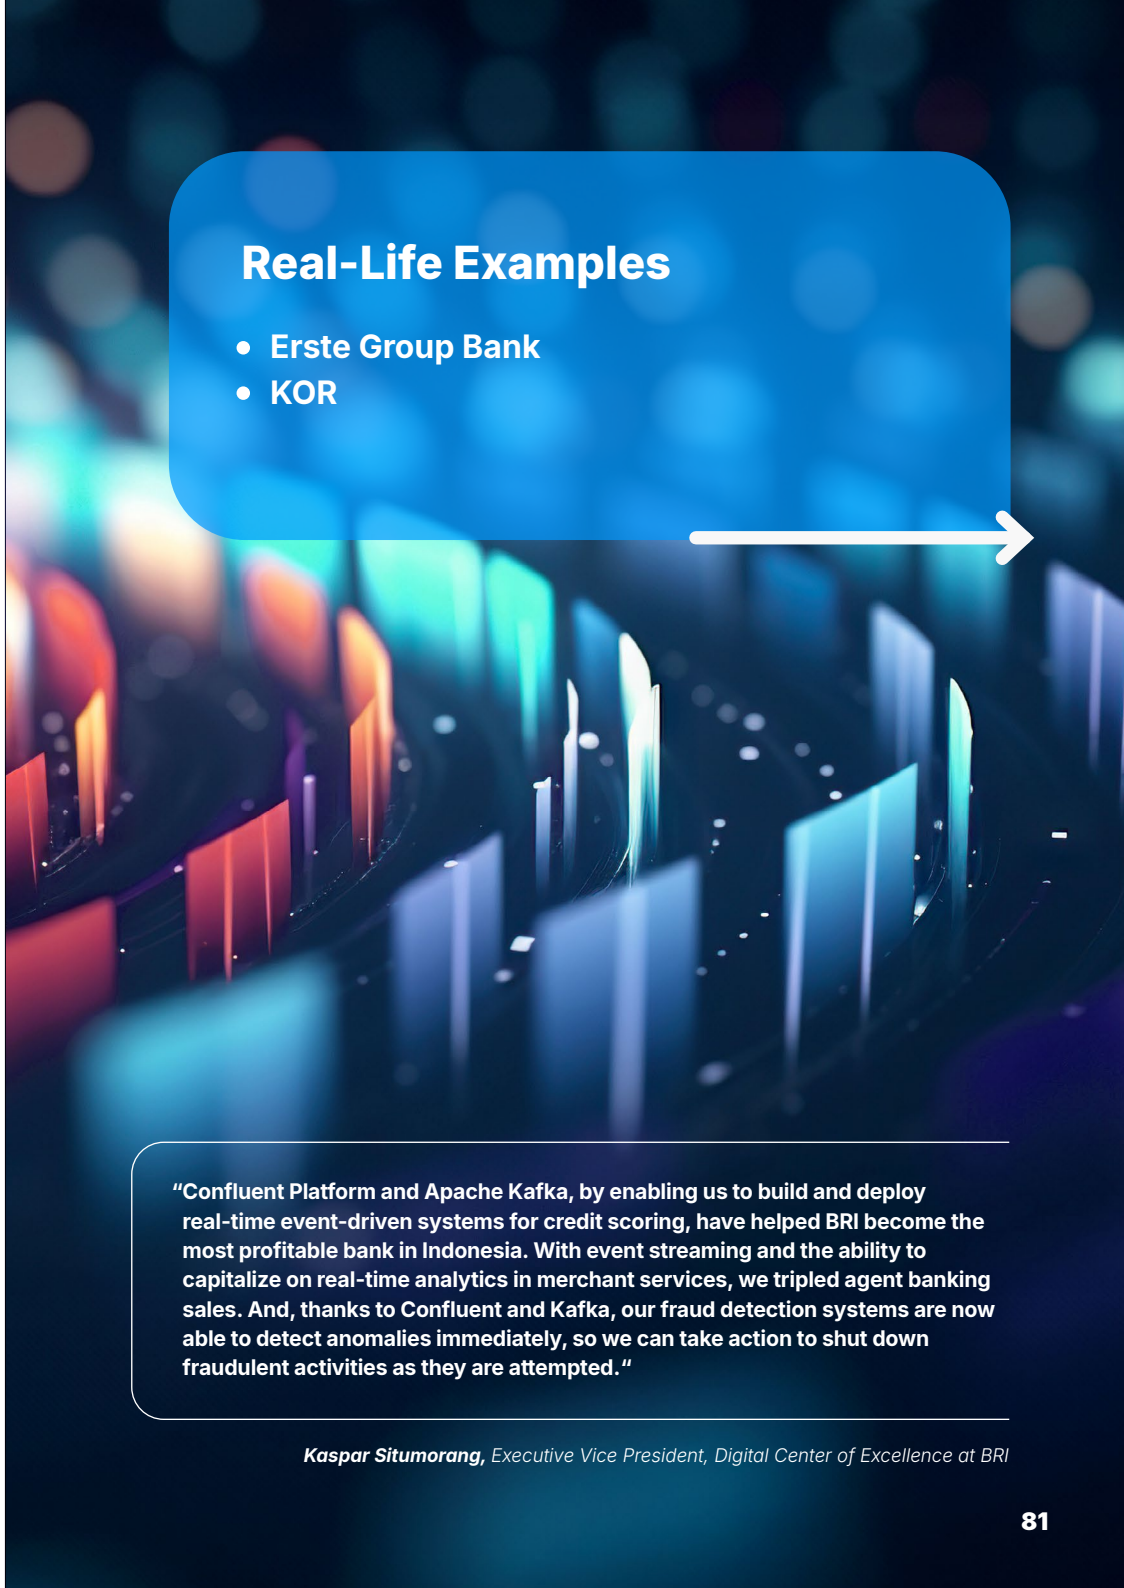

Real-Life Examples

- Erste Group Bank
- KOR

"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. With event streaming and the ability to capitalize on real-time analytics in merchant services, we tripled agent banking sales. And, thanks to Confluent and Kafka, our fraud detection systems are now able to detect anomalies immediately, so we can take action to shut down fraudulent activities as they are attempted."

Kaspar Situmorang, *Executive Vice President, Digital Center of Excellence at BRI*

Real-Life Example Erste Group Bank

Use Case(s): Data Products | IT Modernization | Data Sharing | Data Governance

Company Overview

Erste Group Bank, a leading financial services provider in Central and Eastern Europe, faced escalating costs associated with maintaining their mainframe infrastructure. By offloading data from the mainframe to a cloud-native, microservice-based architecture powered by data streaming, the bank achieved substantial cost savings.

Implementation

Erste's mainframe offloading strategy involved the following steps:

1. **Data extraction:** Data was extracted from the mainframe using Confluent.
2. **Data streaming:** The extracted data was streamed in real time to a cloud-native environment.
3. **Microservices integration:** The data was then processed and managed by a suite of microservices, ensuring seamless integration with existing applications and services.

Erste faced the challenge of maintaining and operating expensive mainframe systems while needing to support modern applications and services. To address this challenge, the bank decided to offload workloads from their mainframe systems to a hybrid cloud environment. The implementation involved leveraging data streaming platforms like Confluent built on top of Apache Kafka to ensure data consistency between the legacy mainframe systems and the new cloud applications.

The data streaming platform facilitated the real-time integration of data from the mainframe systems with the cloud environment. This integration enabled the bank to process and analyze data in real time, ensuring that cloud applications had access to the most up-to-date information. The platform also supported the development of APIs that allowed modern applications to interact seamlessly with the legacy systems.

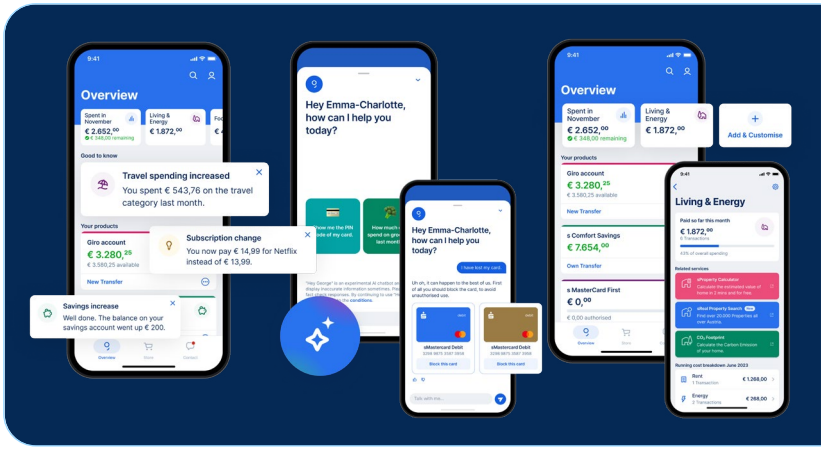

Banking App George by Erste Group Bank AG

Business Value

- **Cost reduction:** The transition allowed the bank to reduce the number of reads on the mainframe, significantly cutting down on fixed infrastructure costs.
- **Compliance and innovation:** The bank maintained compliance with regulatory requirements and business logic while enabling the creation of new applications using the same event-based architecture.
- **Operational efficiency:** The new architecture enhanced the agility of IT operations, allowing for faster deployment of new services and ensuring data consistency across mainframe, databases, and customer-facing applications.

> **"We currently have 214 different projects on our production cluster, 1691 topics, with approx. 50,000 msg/sec incoming at around 75MB/sec, outgoing approx. 100MB/ sec—this last metric is our best indicator of the rising adoption as a data sharing platform: data is read more often than written, i.e., data is produced for multiple consumers."**

Mathias Frey, *Chief Engineering Technology Officer, Erste Group Bank*

Real-Life Example KOR

Use Case(s): IT Modernization | Observability | Data Governance

Company Overview

KOR, a leading trade data repository and reporting service provider for the global derivatives markets, has successfully implemented data streaming technologies to enhance its regulatory reporting capabilities.
By Confluent, KOR has created a robust and efficient regulatory reporting environment that supports real-time data processing and analytics.

Implementation

KOR's regulatory reporting platform continuously streams data from client transaction and lifecycle data, 3rd party trade activity and reference data to its cloud-based platform.
Advanced algorithms process the data in real time to ensure accuracy, identify discrepancies, and generate regulatory reports. Real-time data validation with analytics and support of complex scenario processing set KOR apart from competitors using traditional methods.

Real-time dashboards at KOR

Dashboards and alerts allow operations and compliance to monitor data streams, issues, and rejections to ensure timely and accurate reporting. The data streaming platform also supports ongoing maintenance and updates, ensuring that KOR can focus on innovation and customer service.

Business Value

- **Transparency:** Real-time data validation and transparency improve the accuracy and reliability of regulatory reports.
- **Enhanced compliance:** Continuous monitoring and real-time analytics ensure compliance with regulatory requirements.
- **Operational efficiency:** Handling of complex trading and reporting scenarios along with their remediation actions.
- **Scalability and flexibility:** Cloud-based data streaming platform provides the scalability and flexibility needed to handle large volumes of data and support business growth.
- **Reduced risk:** Real-time data processing and advanced analytics reduce the risk of non-compliance and associated penalties.
- **Improved customer trust:** Enhanced regulatory reporting capabilities improve customer trust and satisfaction.

"At KOR we have a very specific problem that we are trying to solve, which is collecting trading information for regulators. And we decided to do it differently to the way that most people are doing it. Where others would be using Kafka merely as a message bus and use a data lake as the primary data store, we decided to use Kafka as our system of record. We are building our system to store petabytes in Confluent Cloud and maintain the flexibility to materialize projections based on that. So it's a long retention use case."

Andreas Evers, Chief Technology Officer, KOR

Conclusion

The financial services industry is experiencing a profound transformation, with data streaming, AI, and cloud-native solutions driving innovation across customer experiences, business models, and operational efficiencies. Key trends are reshaping the landscape: open banking ecosystems, digital-only banks, and embedded finance solutions are creating new opportunities for revenue and collaboration. Customers now expect mobile-first experiences, instant payment processing, real-time fraud alerts, and virtual financial assistants, requiring financial institutions to operate at the cutting edge of technology.

Real-world examples illustrate the power of these trends. Erste Group Bank modernized its operations by offloading mainframe workloads to a hybrid cloud environment, achieving significant cost savings and operational agility. KOR has transformed regulatory reporting by integrating real-time data validation and AI-driven analytics into its cloud-based platform, ensuring transparency, compliance, and operational efficiency.

Data streaming also optimizes internal efficiencies through automated compliance monitoring, AI-driven risk management, and predictive analytics for market trends. Fraud detection systems built on data streaming platforms provide real-time monitoring, anomaly detection, and operational efficiency, minimizing risks and protecting institutional integrity.

These advancements demonstrate the critical role of data streaming in addressing industry challenges. By adopting scalable, resilient, and innovative architectures, financial institutions can reduce costs, maintain compliance, and enhance customer satisfaction. Transparency, scalability, and flexibility enable institutions to handle vast data volumes while driving rapid innovation and operational excellence.

The convergence of these trends highlights a broader imperative: financial institutions must leverage technologies like data streaming, AI, and hybrid cloud architectures to remain competitive. They can redefine their market leadership by embracing new business models, creating real-time customer experiences, and optimizing internal processes. Institutions like Erste Group Bank, and KOR are already exemplifying what's possible, showing that the future of financial services is built on a foundation of real-time insights, operational agility, and relentless innovation.

3.2 Manufacturing and Automotive

Few industries integrate the physical and digital worlds as effectively as industrial manufacturing and the automotive sector. The potential of digitalization is significant. Industry representatives have been discussing "Industry 4.0" for over a decade which now comes to life. And the trends keep on coming: Think of the adoption of green manufacturing practices, the implementation of digital twins, and the enhancement of cybersecurity measures. Green manufacturing focuses on reducing environmental impact through sustainable practices and energy-efficient technologies.

Digital twins, which are virtual replicas of physical assets, enable real-time monitoring and optimization of manufacturing processes. Cybersecurity is becoming increasingly important as manufacturers adopt more connected and automated systems, making them vulnerable to cyber threats.

The integration of IoT devices and advanced analytics is revolutionizing the way manufacturers monitor and maintain their equipment on the shop floor. Predictive maintenance, often powered by AI, allows manufacturers to predict equipment failures before they occur, reducing downtime and improving overall equipment effectiveness (OEE). Supply chain management is also being transformed through the use of data pipelines and IT modernization, providing end-to-end visibility and enabling more efficient operations. Track-and-trace solutions and real-time locating system (RTLS) leveraging IoT and advanced applications offer real-time insights into the entire supply chain, enhancing efficiency and transparency.

"We chose Confluent Cloud not only because of the best cost benefit ratio but also because it fits the best to our needs, which were the data store capability … near real-time data provisioning … and easy connectivity."

Stefan Baer, *IT Architect and Solution Integrator, Siemens*

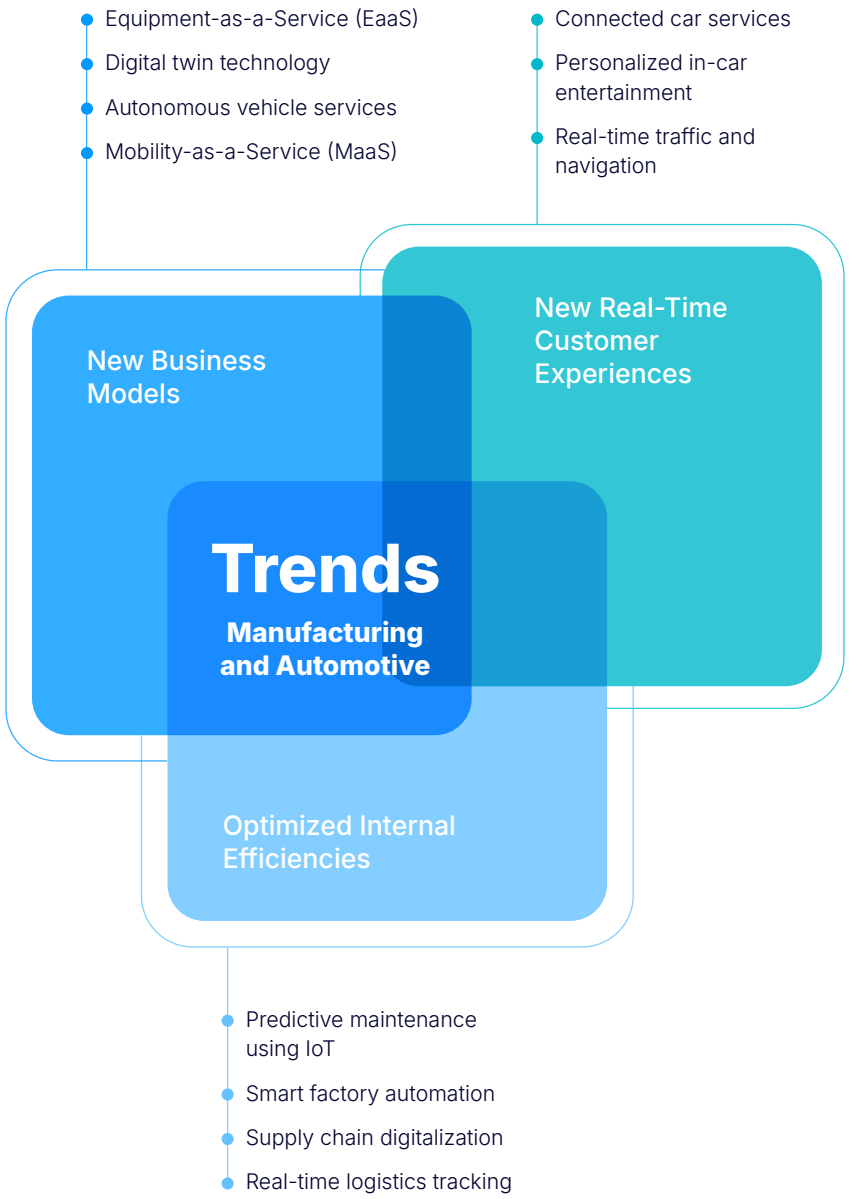

Scenario 1:
Condition Monitoring and Predictive Maintenance

Condition monitoring and predictive maintenance involve continuously tracking shop floor equipment to predict and prevent failures. IoT sensors and advanced analytics collect real-time data on equipment performance, identify patterns, and predict when maintenance is needed.

If this data is then streamed and analyzed in real time on a central platform, it can be used to detect anomalies and automatically alert the maintenance team for them to address the issue before it leads to a breakdown. This proactive approach minimizes downtime, lowers maintenance costs, and enhances OEE.

Business Value

- **Minimized unplanned downtime:** By predicting potential equipment failures before they occur, manufacturers can schedule maintenance proactively, reducing unexpected breakdowns and associated downtime.
- **Reduced maintenance costs:** Transitioning from reactive to predictive maintenance allows for more efficient use of resources, reducing the frequency and cost of maintenance activities.
- **Improved overall equipment effectiveness (OEE):** Continuous monitoring and timely maintenance ensure that equipment operates at optimal efficiency, enhancing productivity and maximizing output.
- **Extended equipment lifespan:** Regular monitoring and maintenance based on actual equipment condition help in extending the lifespan of machinery, delaying the need for costly replacements.
- **Enhanced safety and compliance:** Predictive maintenance ensures that equipment is always in safe working condition, reducing the risk of accidents and ensuring compliance with safety regulations.
- **Increased production quality:** Well-maintained equipment operates more reliably, leading to consistent production quality and reducing the likelihood of defects.
- **Sustainability goals:** By optimizing equipment performance and reducing waste, predictive maintenance contributes to sustainability initiatives and reduces the evironmental impact.

Scenario 2:
Supply Chain

Supply chain management involves coordinating and optimizing various processes like transportation, inventory management, and order fulfillment. Data pipelines and IT modernization enable end-to-end visibility across the entire supply chain for more efficient operations and better decision-making. A truly decoupled, event-driven choreography brings more flexibility and scalability compared to a central orchestration.

Business Value

- **Enhanced visibility and transparency:** Achieve real-time, end-to-end visibility across the entire supply chain, allowing stakeholders to monitor and track the movement of goods and materials seamlessly.
- **Improved decision-making:** Access to real-time data enables more informed and timely decision-making, reducing the risk of errors and enhancing strategic planning.
- **Increased operational efficiency:** Streamlined processes and automated workflows reduce manual intervention, leading to faster order fulfillment and reduced operational costs.
- **Reduced lead times:** Real-time data processing and analytics help in predicting demand and optimizing inventory levels, thereby reducing lead times and improving customer satisfaction.
- **Cost reduction:** By optimizing transportation routes and inventory levels, companies can significantly reduce logistics and warehousing costs.
- **Enhanced collaboration:** Improved data sharing and communication across different departments and partners foster better collaboration and alignment of supply chain activities.
- **Sustainability and compliance:** Efficient resource management and reduced waste contribute to sustainability goals, while real-time data ensures compliance with regulatory requirements.
- **Customer satisfaction:** Faster and more reliable order fulfillment enhances the overall customer experience, leading to increased loyalty and repeat business.

Scenario 3:
Track-and-Trace

Track-and-trace solutions involve the real-time monitoring of tools, equipment, and products throughout the supply chain. By leveraging IoT devices and advanced applications, manufacturers can gain visibility into the location and status of their assets, ensuring transparency and accountability. This real-time visibility helps to prevent losses, improve inventory management, and enhance overall supply chain efficiency.

Business Value

- **Enhanced visibility and transparency:** Real-time tracking of tools, equipment, and products provides complete visibility across the supply chain, ensuring transparency and accountability at every stage.
- **Improved inventory management:** Accurate, real-time data on asset location and status helps optimize inventory levels, reducing overstock and stockouts, and improving inventory turnover.
- **Loss prevention:** Continuous monitoring helps prevent losses due to theft, misplacement, or damage by providing timely alerts and insights into asset movements.
- **Increased supply chain efficiency:** Streamlined tracking processes reduce delays and inefficiencies, leading to faster order fulfillment and improved supply chain responsiveness.
- **Better customer satisfaction:** Real-time updates on product location and status enhance customer communication and satisfaction by providing accurate delivery estimates and reducing uncertainty.
- **Regulatory compliance:** Track-and-trace solutions ensure compliance with industry regulations and standards by maintaining accurate records of product movements and status.

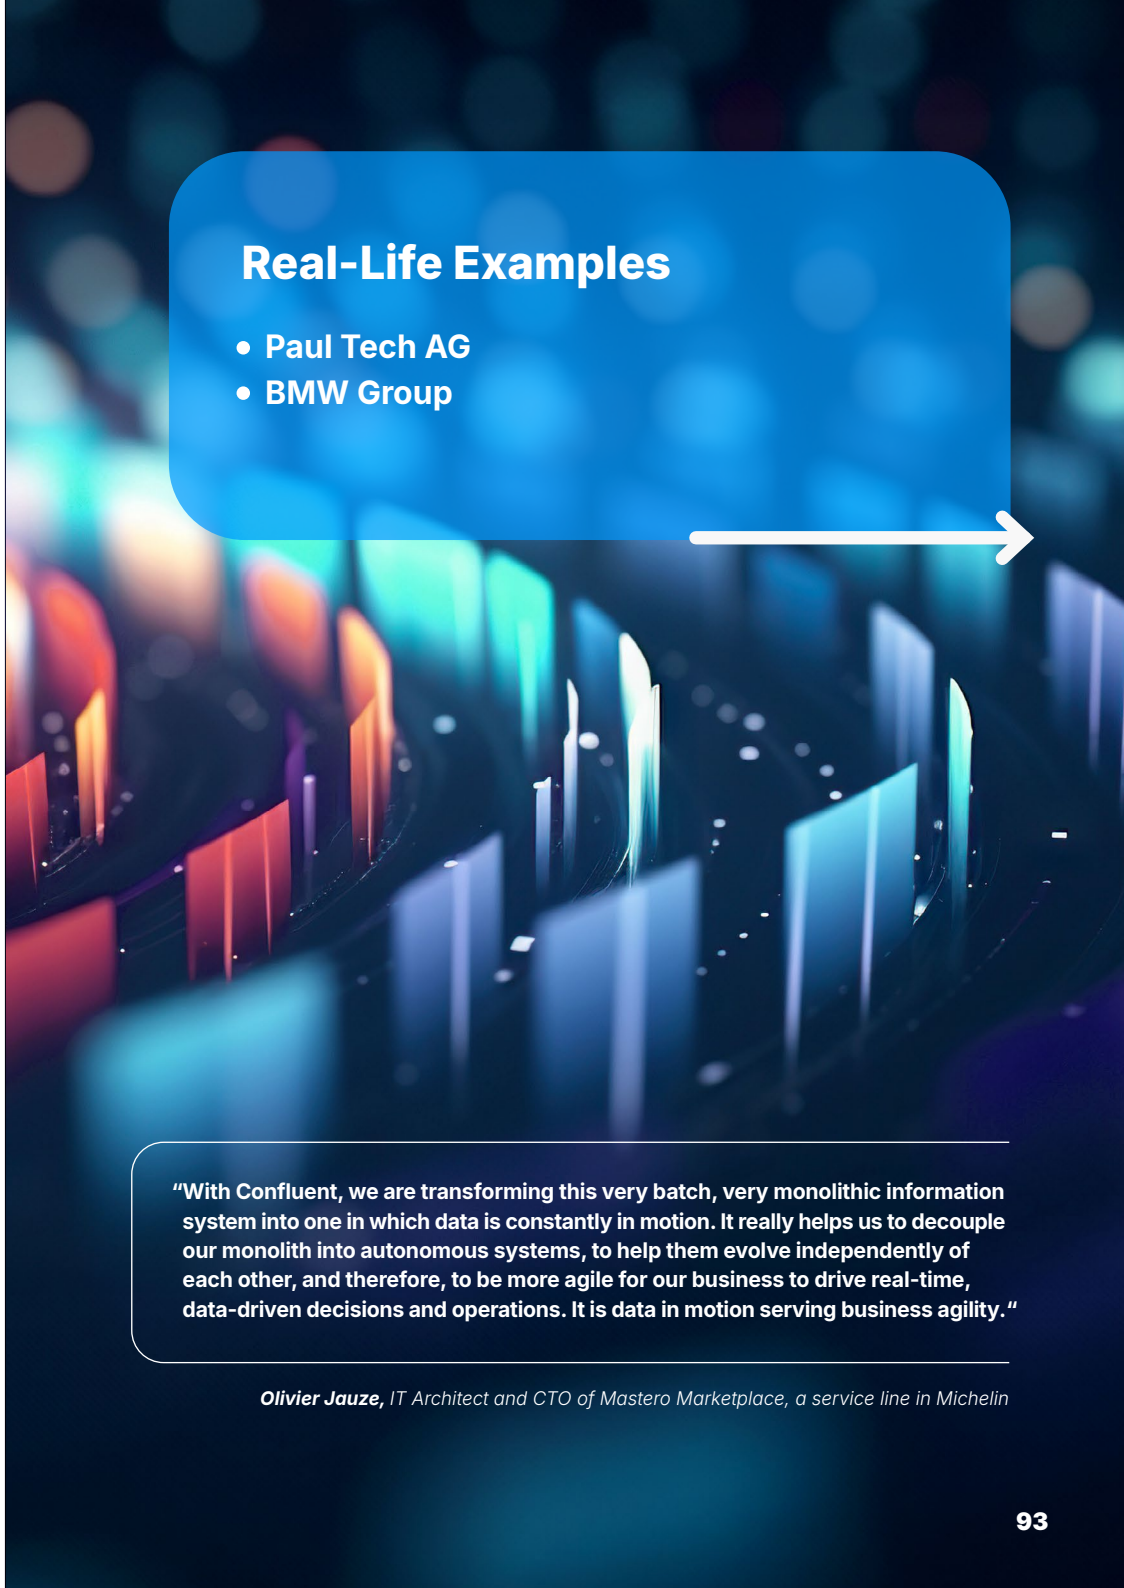

Real-Life Examples

- **Paul Tech AG**
- **BMW Group**

"With Confluent, we are transforming this very batch, very monolithic information system into one in which data is constantly in motion. It really helps us to decouple our monolith into autonomous systems, to help them evolve independently of each other, and therefore, to be more agile for our business to drive real-time, data-driven decisions and operations. It is data in motion serving business agility."

Olivier Jauze, *IT Architect and CTO of Mastero Marketplace, a service line in Michelin*

Real-Life Example PAUL Tech AG

Use Case(s): Hybrid and Multi-Cloud | IoT | ESG

Company Overview

PAUL Tech AG, a pioneering force in the realm of smart building technology and committed to creating a sustainable future by optimizing energy consumption in buildings. With a focus on reducing carbon emissions, PAUL Tech AG has positioned itself as a leader in the industry by leveraging cutting-edge technology to drive significant environmental impact. In 2023 alone, the company is projected to save approximately 21,500 tons of CO_2, a feat equivalent to the carbon absorption capacity of 10,000 square meters of forest. This achievement underscores the company's dedication to sustainability and its innovative approach to energy management.

Implementation

In today's world, where energy conservation and sustainability are paramount, PAUL Tech AG identified a critical challenge: the need to reduce energy consumption in buildings by up to 40% without compromising the comfort of occupants. Traditional methods of energy optimization often require significant physical modifications to existing infrastructure, which can be costly and disruptive. PAUL Tech AG sought a solution that allows for seamless integration with existing systems to leverage data-driven insights to enhance efficiency and sustainability.

To address this challenge, PAUL Tech AG turned to data streaming technology with Confluent. The company implemented an adaptive hydraulic balancing system in drinking water and heating systems, powered by artificial intelligence (AI) to analyze and optimize the current state of the entire system.
This AI-driven approach considers various factors such as consumption behavior, vacancy, weather conditions, sun orientation, and time of day or season, combining billions of data points from other buildings to deliver optimal results.

The integration process involved connecting IoT devices via MQTT from the edge to Confluent Cloud, facilitating the flow of IoT data into the data streaming platform. The right configuration ensured that data from the same edge node was stored together, maintaining the true sequence of events for subsequent processing. Once the data is ingested, stream processing applications leveraged PAUL Tech AG's machine learning models to optimize actuator values. These optimized values were then sent back to the IoT devices.

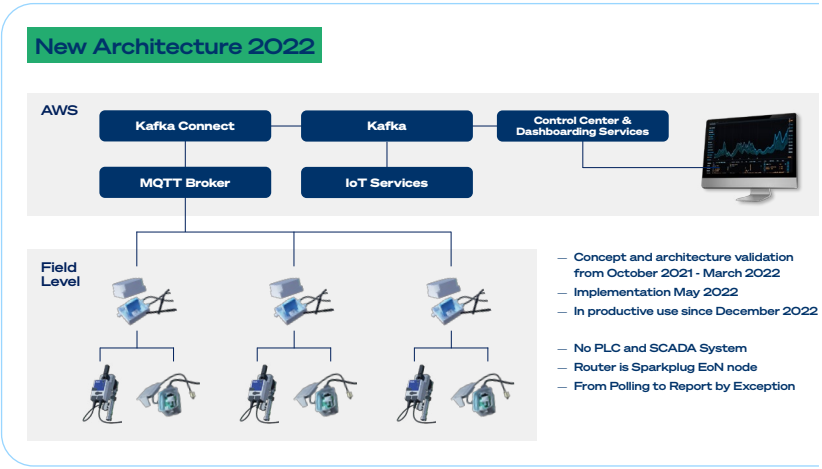

Business Value

- **Significant Energy Savings:**
  Optimizing energy consumption through real-time data processing and machine learning enabled PAUL Tech AG to achieve up to 40% energy savings in buildings.

- **Enhanced ESG and Sustainability:**
  Reduced energy consumption improved ESG ratings, reinforcing PAUL Tech AG's commitment to sustainability and responsible corporate practices.

- **Lower Operational Costs and Financial Advantage:**
  Energy efficiency and improved ESG ratings contributed to reduced operational costs and favorable refinancing rates.

- **Future-Proofed Infrastructure:**
  Data-driven precision in heating system sizing and adaptability for alternative heat sources prepared buildings for future energy innovations.

Real-Life Example BMW Group

Use Case(s): Hybrid and Multi-Cloud | AI | IoT

Company Overview

Founded as Bayerische Motoren Werke GmbH in 1916, today, the BMW Group is the world's leading premium manufacturer of automobiles and motorcycles. The Group operates 31 plants in 15 countries, with a global sales network in 140 countries

Implementation

Inside BMW Group's state-of-the-art plants, data is collected from all the critical minutiae of manufacturing, including machine temperatures, machine performance, and location of parts on the assembly line. Data is also continuously produced and consumed by sales and the research and development (R&D) operation.

Making the abundance of real-time Internet of Things (IoT) data available to all teams that want to put it to use is a critical role for BMW Group's Integration Platforms organization to drive innovation through use cases such as enabling operational applications that streamline factory production and providing real-time data to the R&D team as they create and test new products.

Data flows continuously through Confluent via a combination of commercial solutions and custom code, using streaming ETL—the process of continuously transforming, filtering, and enriching data as it moves. As Confluent's Infinite Storage persists the incoming events, BMW Group can also replay historical data later to enable additional use cases. Any asset that produces data is connected to Kafka, with multiple different application types using the same data to trigger different, specific processes.

BMW Group operates a complex infrastructure, with a large number of SAP systems, a lot of legacy software, plus third-party developed connectors and tools. The company now has dozens of clusters, hundreds of connected applications, and over one billion events produced per day—all with Confluent functioning as the central nervous system for this continuous river of critical business data. Highly scalable teams at BMW Group are onboarding multiple new applications and use cases every week.

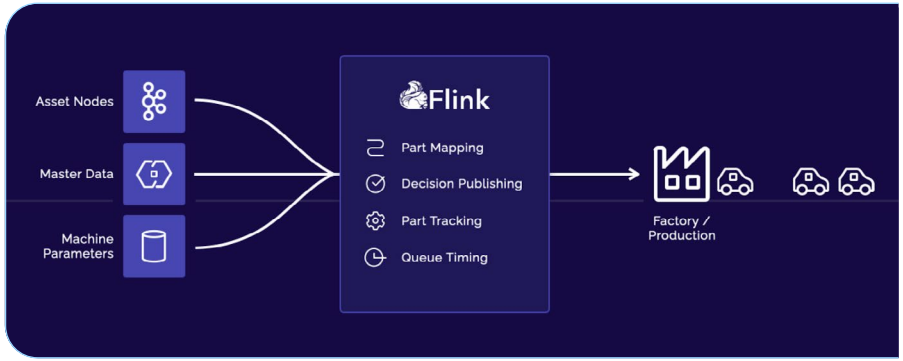

"BMW Group's successful implementation of Confluent has already resulted in increased speed, flexibility, cost savings, and simplicity. Yet the company is ambitious and keen to expand the use of the data processing platform."

BMW Group

Business Value

The implementation of data streaming technology has provided significant business value for BMW Group:

- **Enhanced operational efficiency:** Real-time data processing and analytics enable BMW Group to optimize production processes, reduce downtime, and improve overall equipment effectiveness (OEE).
- **Cost and risk reduction:** By leveraging fully managed cloud service for data processing and eliminating the need for manual infrastructure management, BMW Group has reduced operational costs and risks.
- **Improved innovation and flexibility:** The scalable nature of Confluent allows BMW Group to accelerate project timelines and increase opportunities for innovation. The ability to process vast amounts of data in real time supports continuous IoT innovation and enhances customer engagement.
- **Ubiquitous access to real-time data:** Confluent provides BMW Group with a powerful platform for accessing, storing, and managing data. Self-service access to data products enables teams to make informed decisions quickly, driving business growth, efficiency and faster time to market.

Conclusion

The manufacturing and automotive industries are at the forefront of digital transformation, driven by business models such as Equipment-as-a-Service (EaaS) and Mobility-as-a-Service (MaaS), alongside innovative technologies like digital twin systems that enhance operational efficiency and product development.

These new business models and technology advancements are revolutionizing the shop floor, while real-time customer experiences redefine the buying process and aftersales. IoT-driven automation, smart factory workflows, and supply chain digitalization deliver optimized efficiency and sustainability across the value chain.

BMW Group embodies this transformation by harnessing Confluent to collect, process, and distribute real-time IoT data across 31 global plants. Streaming ETL enables the continuous enrichment of data, powering use cases from factory automation to real-time R&D insights. Processing over a billion events daily, BMW Group achieves significant cost reductions, enhanced operational efficiency, and accelerated innovation, solidifying its leadership in sustainable manufacturing.

PAUL Tech AG illustrates the sustainability potential of data streaming by optimizing energy consumption in buildings through IoT and AI. Its real-time analytics platform achieves up to 40% energy savings, reducing operational costs while enhancing ESG ratings. These results highlight how Confluent empowers organizations to meet environmental goals without compromising operational performance.

Condition monitoring and predictive maintenance further demonstrate the value of IoT and AI in minimizing downtime maximizing asset performance, and improving the OEE. Smart factory automation enhances production precision and flexibility, while supply chain digitalization improves logistics tracking, reducing lead times and fostering sustainability.

These examples underscore the transformative power of data streaming, IoT, and AI in driving efficiency, innovation, and sustainability. Companies like BMW Group, and PAUL Tech AG showcase how continuous real-time data processing enables scalable, event-driven architectures that redefine operations and customer experiences. As industries embrace these technologies, they position themselves for a future where agility, sustainability, and customer-centricity are paramount.

3.3 Retail and E-Commerce

The retail industry is experiencing a wave of digital transformation driven by the need for omnichannel experiences and unified commerce, automation through AI, retail media networks, and the influence of social media. Unified commerce integrates various shopping channels, providing a seamless customer experience across online and offline platforms. Automation through AI enables retailers to optimize operations, personalize customer interactions, and improve decision-making. Retail media networks allow companies to monetize their digital platforms by offering advertising opportunities to brands. Social media plays a crucial role in shaping consumer behavior and driving engagement through interactive and personalized content.

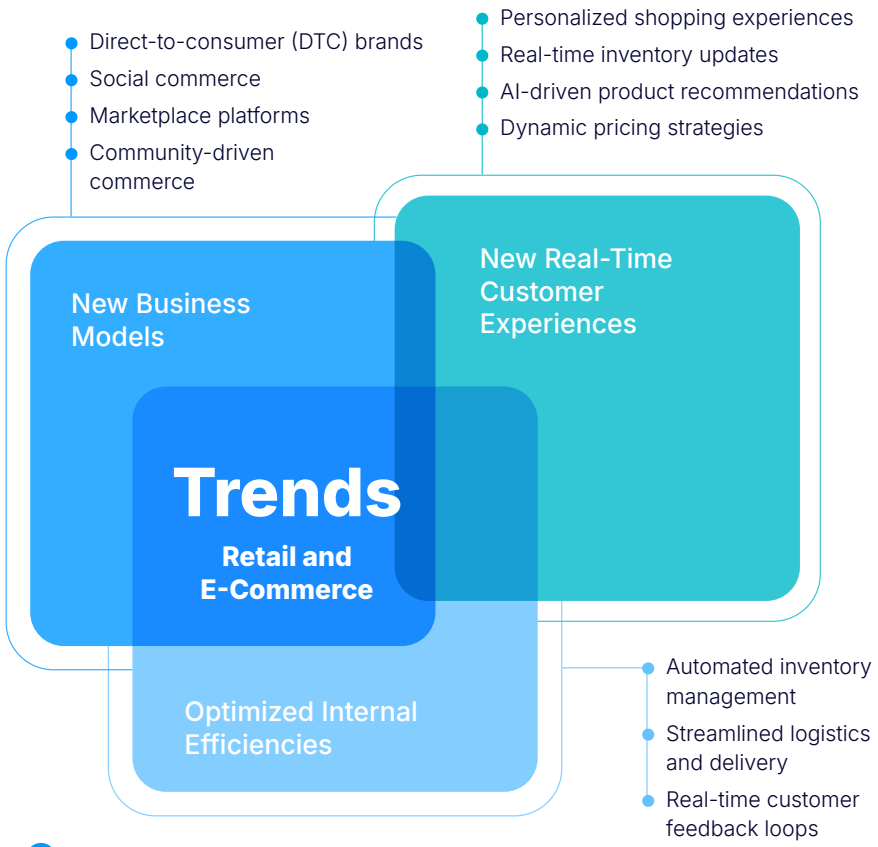

Scenario 1:
Inventory Management

Inventory management ensures optimal availability while reducing costs through tracking and replenishment. With data pipelines and IT modernization, retailers can use real-time data from point-of-sale systems, supplier databases, and demand patterns to enhance forecasting and replenishment processes.

Real-time data streaming helps retailers to reduce excess inventory, minimize stockouts, and improve overall operational efficiency, resulting in significant cost reductions.

Business Value

- **Improved forecasting accuracy:** Real-time data from various sources enhances forecasting models, allowing retailers to predict demand more accurately and adjust inventory levels accordingly.
- **Optimal stock levels:** By maintaining the right balance of inventory, retailers can reduce excess stock and minimize stockouts, ensuring product availability without overstocking.
- **Reduced inventory costs:** Efficient inventory management reduces holding costs and minimizes the financial impact of unsold stock, leading to significant cost savings.
- **Enhanced replenishment processes:** Real-time data processing enables more efficient and timely replenishment, ensuring that products are available when and where they are needed.
- **Increased operational efficiency:** Streamlined inventory processes reduce manual intervention and errors, improving overall operational efficiency and productivity.
- **Better supplier collaboration:** Real-time data sharing with suppliers enhances collaboration and coordination, leading to more efficient supply chain operations.
- **Improved customer experience:** By ensuring products are consistently available and reducing stockouts, retailers can enhance customer satisfaction and loyalty, leading to repeat business and positive brand perception.

Scenario 2:
Loyalty Platforms

Retailers implement a loyalty platform powered by data streaming to significantly enhance customer engagement and drive incremental revenue. Data streaming technologies play a crucial role in this process by continuously collecting and processing real-time data from various sources, such as transaction history, membership activity, and feedback channels. This enables companies to create dynamic and responsive loyalty programs that offer points, rewards, and exclusive benefits tailored to individual customer behaviors and preferences. With real-time transaction monitoring, retailers can instantly update loyalty points and rewards balances, ensuring customers receive timely recognition for their loyalty. Additionally, the integration of instant feedback allows retailers to quickly adapt loyalty program features to better meet customer expectations, enhancing overall satisfaction.

Business Value

- **Enhanced customer loyalty:** Real-time data processing allows retailers to offer timely and relevant rewards, fostering a sense of appreciation and loyalty among customers.
- **Increased customer engagement:** Dynamic loyalty programs encourage active participation, as customers are motivated to earn and redeem rewards, leading to higher engagement levels.
- **Improved brand perception:** Offering exclusive benefits and personalized rewards enhances the brand's image, making it more attractive to both existing and potential customers.
- **Data-driven insights:** Real-time data collection and analysis provide valuable insights into customer behavior and preferences, enabling retailers to refine their loyalty strategies and offerings.
- **Competitive advantage:** A well-executed loyalty platform differentiates the brand in a crowded market, attracting more customers and retaining existing ones.
- **Drive additional revenue:** Effective loyalty programs can drive additional spend and can be monetized through business partnerships.

Scenario 3:
Personalization

Data streaming technologies empower retailers to offer personalized shopping experiences to their customers. By processing real-time data from diverse sources such as purchase history, browsing behavior, and demographic information, companies can customize product recommendations, marketing messages, and shopping experiences to align with individual preferences. This capability in real-time data processing enables retailers to boost customer satisfaction, elevate conversion rates, and foster stronger customer relationships.

Business Value

- **Personalized customer experience:** Real-time data processing enables retailers to offer tailored product recommendations and marketing messages, creating a more engaging and relevant shopping experience.
- **Increased conversion rates:** Personalized offers and recommendations drive higher conversion rates, as customers are more likely to purchase products that align with their interests and needs.
- **Enhanced customer satisfaction:** By delivering a shopping experience that resonates with individual preferences, retailers can improve customer satisfaction and encourage repeat visits.
- **Stronger customer relationships:** Personalization fosters a deeper connection between the brand and its customers, leading to long-term relationships and increased loyalty.
- **Optimized marketing strategies:** Data-driven insights allow retailers to refine their marketing strategies, ensuring that campaigns are targeted and effective in reaching the right audience.

"Real-time data is important to AO.com because we deliver real-time personalized banners to our users as they navigate our site. Confluent unlocks these types of use cases that are very important to our business."

Amelia McVennon, *Developer, AO.com*

Scenario 4:
Live Commerce

Live commerce involves the use of social commerce for live auctions and interactive selling, providing a dynamic and engaging shopping experience. By leveraging data pipelines, log aggregation, and advanced applications, retailers can enable real-time interactions between sellers and buyers, driving engagement and sales. This real-time data processing capability helps retailers to innovate and create new revenue streams through interactive and personalized shopping experiences.

Business Value

- **Dynamic and engaging shopping experience:** Real-time interactions between sellers and buyers create a lively and interactive shopping environment, enhancing customer engagement and enjoyment.
- **Increased sales and conversion rates:** The immediacy and interactivity of live commerce encourage impulse purchases and drive higher conversion rates, boosting sales.
- **New revenue streams:** By innovating with live auctions and interactive selling, retailers can tap into new revenue opportunities and diversify their sales channels.
- **Improved brand visibility and reach:** Live commerce events can attract a wide audience, increasing brand visibility and expanding reach to new customer segments.

Real-Life Examples

- **Migros**

"Our ability to preserve the timing of our ERP program launch saved millions of dollars in consulting and resources that would have been needed to keep our old system going. We would never have been able to keep that timing if we did not have the new, flexible event-driven architecture enabled by Confluent."

Beth Zuke, *North America Technology Director at Amway*

Real-Life Example Migros

Use Case(s): Data Products | Observability | Event-Driven Microservices | IoT

Company Overview

Migros is one of Switzerland's largest retail companies, supermarket chains, and employers, known for its extensive network and diverse range of products and services. As a prominent player in the Swiss retail market, Migros is committed to innovation and efficiency, particulary in its supply chain operations. The company has embraced data streaming technologies to enhance logistics and overall supply chain modernization.

Implementation

Migros faced the challenge of optimizing its supply chain to achieve a global optimum rather than settling for local optima. The company needed to improve its logistics and track-and-trace capabilities to provide real-time transportation information and enhance decision-making processes. The existing systems were prone to delays and inefficiencies, with challenges such as non-standardized data ingestion, resource-intensive operations, and static workflows that limited agility.

To address these challenges, Migros implemented a single data streaming pipeline using Confluent and IoT connectors. This solution enabled real-time visualization and processing of logistics and transportation information, allowing Migros to forecast truck arrival times and reschedule truck tours effectively. By leveraging Confluent, Migros was able to distribute master data across multiple systems and ensure seamless integration and consistent synchronization of information.

Migros' journey with data streaming began with the analytics team tracking customer events across different touchpoints. Over time, the company expanded its use of Confluent to include customer master data management, product and store information updates, and synchronization with Migros Online. The implementation of stream processing applications further enhanced Migros' ability to process data in real time and support advanced analytics. One of the next projects on Migros' journey is the implementation of track-and-trace capabilities for a clear and up-to-date view of inventory levels, shipment statuses, and delivery timelines.

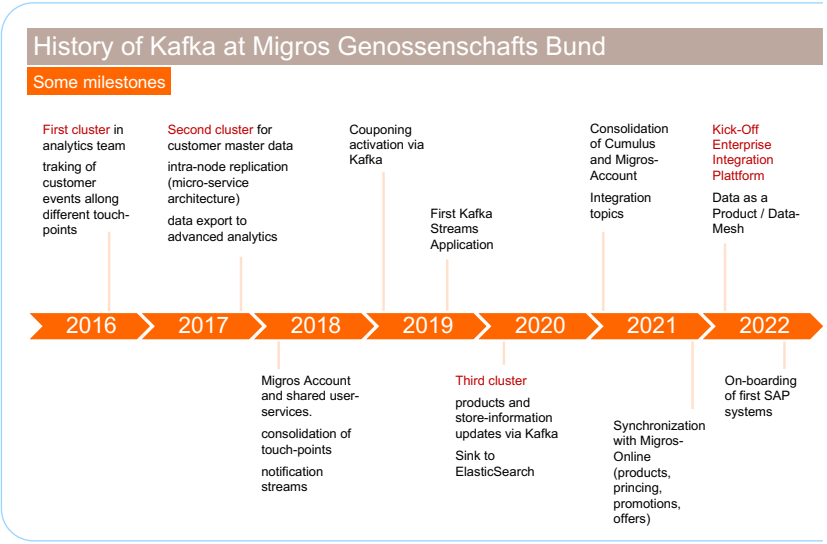

History of Kafka at Migros Genossenschafts Bund

Business Value

- **Optimized supply chain efficiency:** A single streaming pipeline enabled Migros to achieve a global optimum, enhancing supply chain efficiency and reducing operational costs.
- **Improved decision-making with real-time insights:** Real-time transportation and logistics data allowed for prompt responses to events, preventing lost opportunities.
- **Enhanced customer experience:** Consolidating customer touchpoints and integrating key services streamlined operations, providing customers with timely and relevant information.
- **Increased agility and reduced delays:** Transitioning to faster choreography processes improved system scalability, reduced information propagation delays, and increased supply chain agility.

Conclusion

Retail and e-commerce are undergoing a significant transformation, driven by trends such as direct-to-consumer (DTC) brands, social commerce, and marketplace platforms. Unified commerce integrates shopping channels for a seamless experience, while automation and AI optimize inventory management, logistics, and customer personalization. Real-time customer experiences, such as AI-driven product recommendations, dynamic pricing, and real-time inventory updates, are reshaping how retailers engage with consumers.

Migros exemplifies the power of data streaming in optimizing supply chains. Migros created a single data streaming pipeline to achieve real-time logistics tracking and supply chain visibility. This transformation reduced operational costs, increased agility, and enhanced customer satisfaction by ensuring accurate delivery timelines and synchronized data across systems.

Retailers are also capitalizing on live commerce, combining social commerce with interactive live selling to create dynamic customer experiences. By leveraging data pipelines, brands can enhance engagement and drive new revenue streams through real-time interactions between buyers and sellers.

These advancements highlight the critical role of data streaming in modern retail. Real-time insights drive personalized shopping experiences, boosting conversion rates and strengthening customer relationships. Simultaneously, operational efficiencies like automated inventory management and real-time feedback loops lower costs and improve supply chain resilience.

Companies like Migros showcase the transformative impact of integrating data streaming into retail operations. By leveraging personalization, agility, and real-time insights, businesses can thrive in a competitive market while adapting to evolving consumer expectations.

These data-driven strategies optimize processes and position retailers for sustainable, long-term growth.

3.4 Telco and Media

The telco and media industry is at the forefront of digital transformation, driven by the rapid adoption of 5G, IoT, and advanced data analytics. The deployment of 5G networks is revolutionizing connectivity, offering ultra-low latency and high-speed data transfer, which is essential for real-time applications and services like multi-device video streaming. IoT is enabling the creation of smart environments, from connected homes to smart cities, providing new opportunities for telco and media companies to offer innovative services.

Personalization is another significant trend, with consumers expecting tailored content and experiences. Advanced analytics and AI/ML are being used to analyze user behavior and preferences, enabling companies o deliver personalized content and advertisements. The impact of interactive gaming is also noteworthy, as it drives demand for high-performance networks and real-time data processing capabilities.

The integration of data streaming platforms is transforming the way telco and media companies manage their operations and deliver services. These platforms enable real-time data processing, providing insights that drive decision-making and enhance customer experiences. As the industry continues to evolve, the adoption of data streaming technologies will be crucial in maintaining competitiveness and driving innovation.

"Data at telco scale is massive. We measure in terms of hundreds of petabytes per day. But in order to start extracting and understanding that data, we partner with companies who understand what the data can mean, how to transact on it, and how to realize the value. And in order for those communications to happen and for us to share data between different clouds, we want to have a broker of information. Confluent is our partner for this."

DISH Wireless

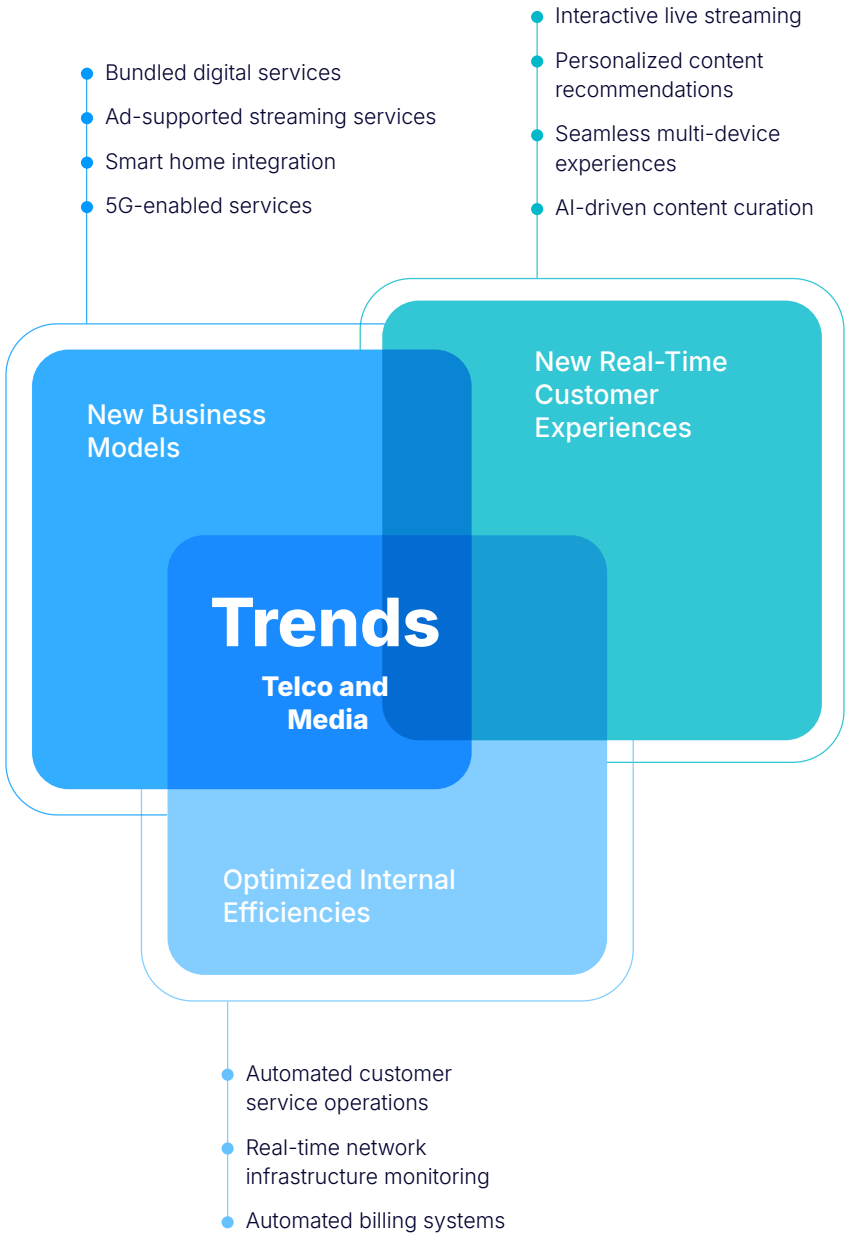

Scenario 1:
Network Monitoring and Alerting

Network monitoring and alerting involve continuously monitoring network performance to detect issues before they impact customers. IT modernization, log aggregation, and multiple clusters offer real-time network insights and help ensure high availability and reliability.

Telecommunication companies can implement network monitoring and alerting systems processing real-time data from various sources, including network devices, log files, and performance metrics. Data streaming is for example used to fully automate site provisioning, service ordering, and service assurance monitoring and reporting for their television networks. This helps telco companies to improve network performance, reduce downtime, and enhance customer satisfaction.

Business Value

- **Real-time network visibility:** Continuous monitoring provides telco companies with real-time insights into network performance, enabling proactive management and optimization.
- **Improved network performance:** By detecting and addressing potential issues promptly, telcos can maintain optimal network performance and ensure high availability and reliability.
- **Reduced downtime:** Real-time alerting allows for quick identification and resolution of network issues, minimizing downtime and service disruptions for customers.
- **Scalability and flexibility:** Data streaming solutions provide the scalability and flexibility needed to adapt to changing network demands and technological advancements.
- **Proactive issue detection:** The ability to detect potential issues before they impact customers allows telcos to take preventive measures, reducing the likelihood of major outages.
- **Enhanced customer satisfaction:** Reliable network performance and reduced service interruptions lead to higher levels of customer satisfaction and loyalty.

Scenario 2:
OTT Services

Over-the-top (OTT) services deliver personalized on-demand content to consumers via the internet, bypassing traditional channels. Using log aggregation, advanced applications, and AI/ML, media companies can analyze user behavior, content metadata, clickstream data, and viewing history in real time. This in turn allows them to provide personalized content recommendations and improve the overall user experience, helping to drive engagement, increase revenue, and enhance customer satisfaction.

Business Value

- **Personalized content recommendations:** Real-time data processing allows media companies to deliver tailored content suggestions based on individual user preferences, enhancing the overall viewing experience.
- **Increased user engagement:** Personalized recommendations and on-demand content keep users engaged and encourage longer viewing sessions, boosting platform usage.
- **Enhanced customer satisfaction:** By providing content that aligns with user interests, OTT services can improve customer satisfaction and foster loyalty.
- **Data-driven insights:** Access to real-time user behavior and content performance data supports more informed decision-making, allowing for strategic content acquisition and marketing efforts.
- **Improved content discovery:** Personalized recommendations help users discover new content, increasing the visibility and consumption of a broader range of offerings.
- **Competitive advantage:** OTT platforms that offer personalized and engaging experiences can differentiate themselves in a crowded market, gaining a competitive edge.

Scenario 3:
Advertisements

Targeted advertisements deliver personalized and relevant ads to consumers based on their preferences and behavior. With advanced applications and AI/ML, media companies can analyze consumer data like behavior, app usage, and demographic information in real time. This way, they are able to deliver ads that are more likely to resonate with their audience.

Data streaming helps companies to optimize ad spend through automation, behavioral analytics, audience targeting, and rewards programs, resulting in improved effectiveness of advertising campaigns and increased ad revenue.

Business Value

- **Enhanced ad relevance:** Real-time data processing allows media companies to deliver personalized and relevant ads based on individual user preferences and behaviors, increasing the likelihood of engagement.
- **Improved campaign effectiveness:** Targeted advertisements are more likely to resonate with the audience, leading to higher conversion rates and improved overall campaign performance.
- **Increased ad revenue:** By delivering more effective ads, media companies can attract higher ad spend from advertisers, boosting ad revenue.
- **Optimized ad spend:** Real-time analytics and automation enable more efficient allocation of ad budgets, ensuring that resources are directed toward high-performing campaigns and audiences.
- **Behavioral analytics:** Access to real-time user behavior data supports more accurate audience segmentation and targeting, enhancing the precision of advertising efforts.
- **Audience targeting:** Leveraging demographic and behavioral data allows for more precise audience targeting, ensuring that ads reach the most relevant users.
- **Real-time engagement:** Delivering ads in real time ensures that they are timely and contextually relevant, increasing the chances of user interaction and engagement.

Conclusion

The telco and media sectors are undergoing profound changes driven by the deployment of 5G networks, the expansion of IoT, and the integration of real-time data analytics. Emerging business models such as bundled digital services, ad-supported streaming platforms, and 5G-enabled smart home integrations are creating new revenue streams and reshaping customer expectations. Real-time customer experiences, including interactive live streaming, personalized content recommendations, and seamless multi-device connectivity, are now essential for engaging audiences. Internally, automated systems for customer service, billing, and network monitoring improve efficiency and scalability.

OTT services showcase the potential of personalized content delivery. By analyzing real-time data from user behavior and preferences, platforms deliver tailored recommendations that enhance customer satisfaction and increase engagement. Targeted advertising, enabled by real-time behavioral analytics, improves ad relevance, boosts campaign effectiveness, and increases ad revenue through precision targeting.

Network monitoring powered by data streaming ensures real-time infrastructure visibility, proactive issue detection, and reduced down-time. These capabilities maintain high network performance and reliability, enhancing customer satisfaction and loyalty. Telco and media companies also leverage high-speed internet and reliable telecommunications infrastructure to support applications like multi-device streaming, smart home integration, and interactive gaming.

These examples highlight the pivotal role of data streaming in telco and media operations. Continuous real-time data processing enables businesses to deliver personalized experiences, optimize internal workflows, and adapt to evolving customer demands. Companies like DISH Wireless demonstrate how integrating technologies like AI, IoT, and real-time analytics drives operational excellence and customer-centric innovation.

As the industry evolves, the adoption of data streaming platforms will be critical for telco and media companies to stay competitive, meet consumer expectations, and unlock new growth opportunities. From improving network reliability to delivering engaging, personalized content, the future of telco and media is being shaped by real-time insights and scalable, data-driven solutions.

3.5 Gaming

Two of the most notable trends the gaming industry has undergone in the past few years is the rise of mobile gaming and Gaming as a Service (GaaS). Mobile gaming continues to dominate the market, with smartphones becoming the primary gaming device for many users. This shift has led to the development of more sophisticated and engaging mobile games, supported by advanced data analytics and AI/ML technologies. GaaS allows players to access games on demand via cloud platforms, eliminating the need for high-end hardware.

Monetization innovation is another critical trend, with game developers exploring new ways to generate revenue through in-game purchases, advertisements, and subscriptions. The usage of data streaming platforms is revolutionizing the gaming industry by enabling real-time data processing and analytics. This capability allows game developers to deliver personalized experiences, optimize game performance, and enhance player engagement.

The importance of real-time data cannot be overstated, as it enables live monitoring, troubleshooting, and the ability to make server-side changes while the game is in progress.

The rise of esports and interactive gaming has further increased the demand for robust and scalable infrastructure to support millions of concurrent players. As the gaming industry continues to evolve, the adoption of data streaming technologies will be crucial in maintaining competitiveness and driving innovation.

"We work closely with the studios to collaborate from the early stages of development. From there, we create our own internal tools and tech to be efficient, responsive and reusable. We enjoy taking on real challenges that serve millions of players every day."

Tom Szymanski, *Team Lead, Data Pipelines at Demonware*

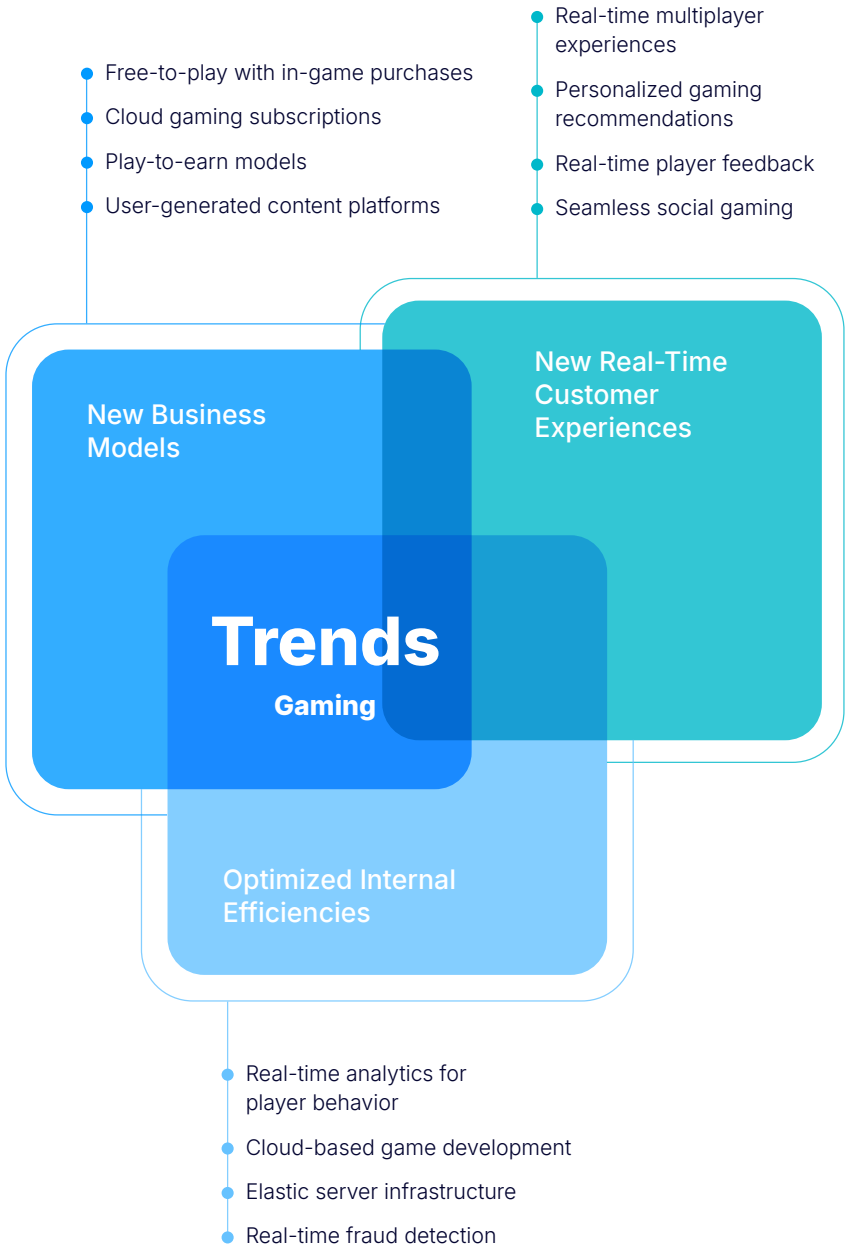

Scenario 1:
Infrastructure Operations

Infrastructure operations in the gaming industry involve the continuous monitoring and management of game servers and networks to ensure optimal performance and availability. By leveraging data pipelines, log aggregation, and multiple clusters, gaming companies can gain real-time visibility into their infrastructure and quickly identify and resolve issues. This real-time data processing capability helps gaming companies to provide 24/7 service health, perform root cause analysis in case of issues, and support cloud-native elasticity for traffic spikes, such as during a new game release or esports tournament.

Data streaming allows gaming companies to serve millions of gamers worldwide without outages, ensuring a seamless gaming experience. Cloud-native infrastructure operations automate site provisioning, service ordering, and service assurance monitoring and reporting. Hence companies can process vast amounts of data in real time, ensuring business continuity without outages.

Business Value

- **24/7 service health:** Continuous monitoring and real-time data processing ensure that game servers and networks maintain optimal performance and availability, providing uninterrupted service to gamers worldwide.
- **Seamless gaming experience:** By preventing outages and minimizing latency, gaming companies can deliver a smooth and enjoyable gaming experience, enhancing player satisfaction and retention.
- **Rapid issue resolution:** Real-time visibility into infrastructure allows for quick identification and resolution of issues, reducing downtime and maintaining service quality.
- **Cloud-native elasticity:** The ability to scale infrastructure dynamically supports traffic spikes during events like new game releases or esports tournaments, ensuring consistent performance under varying loads.
- **Cost efficiency:** Optimized resource allocation and automated processes reduce operational costs, maximizing the return on infrastructure investments.

Scenario 2:
Game Telemetry

Game telemetry involves collecting and analyzing data generated by players during gameplay. Gaining insights into player behavior and preferences through large-scale observability and advanced applications allows gaming companies to deliver personalized experiences and improve game performance.

This real-time data helps gaming companies to make server-side changes while the player is playing the game, such as time-limited events that give rewards or real-time updates to improve the game or align with audience needs.

Business Value

- **Personalized player experiences:** Real-time analysis of player behavior and preferences allows gaming companies to deliver tailored experiences, enhancing player engagement and satisfaction.
- **Dynamic game updates:** The ability to make server-side changes in real time enables the introduction of time-limited events, rewards, and updates that keep the game fresh and exciting for players.
- **Improved game performance:** Continuous monitoring and analysis of game telemetry data help identify performance bottlenecks and optimize game mechanics, ensuring a smooth and enjoyable gaming experience.
- **Increased player petention:** By offering personalized and engaging experiences, gaming companies can increase player retention and foster long-term loyalty.
- **Data-driven game development:** Insights gained from telemetry data inform game design and development, allowing companies to align their offerings with player preferences and market trends.
- **Real-time feedback loop:** Immediate access to player data creates a feedback loop that enables rapid iteration and improvement of game features and content.
- **Enhanced monetization strategies:** Understanding player behavior and preferences allows for the development of targeted monetization strategies, such as in-game purchases and premium content offerings.

Scenario 3:
Monetization Network

Monetization networks in the gaming industry generate revenue via in-game purchases, advertisements, and microtransactions. Leveraging AI/ML and real time data streaming, gaming companies can analyze player behavior and preferences in real-time to deliver targeted advertisements and personalized in-game offers. This helps them optimize ad spend, increase player engagement, and drive revenue growth through upselling and cross-selling opportunities.

Business Value

- **Increased revenue generation:** Real-time data processing enables the delivery of targeted advertisements and personalized in-game offers, maximizing upselling and cross-selling opportunities to drive revenue growth.
- **Optimized ad spend:** By analyzing player behavior and preferences in real time, gaming companies can optimize their advertising strategies, ensuring that ad spend is directed toward the most effective campaigns and audiences.
- **Enhanced player engagement:** Personalized in-game offers and advertisements resonate more with players, increasing engagement and interaction with monetization features.
- **Data-driven monetization strategies:** Access to comprehensive player data supports the development of informed and effective monetization strategies, aligning offers with player preferences and behaviors.
- **Real-time adaptability:** The ability to process and analyze data in real time allows gaming companies to quickly adapt their monetization tactics to changing player behaviors and market conditions.
- **Sustainable business growth:** By maximizing monetization opportunities and driving revenue growth, gaming companies can achieve sustainable business success and long-term profitability.

Conclusion

The gaming industry is at the forefront of digital innovation, driven by trends such as free-to-play models with in-game purchases, cloud gaming subscriptions, and play-to-earn platforms. The rise of mobile gaming and Gaming as a Service (GaaS) has further expanded accessibility, allowing players to enjoy sophisticated experiences on demand. Real-time personalized recommendations, multiplayer gaming, and seamless social experiences are transforming player engagement, while continuous analytics, elastic server infrastructures, and fraud detection ensure operational efficiency.

Game telemetry is another critical application of data streaming, providing continuous insights into player behavior and enabling real-time adjustments. Developers use telemetry data to introduce time-limited events, optimize game mechanics, and deliver personalized rewards, keeping games fresh and engaging. This approach improves player retention, informs data-driven game development, and supports enhanced monetization strategies.

Monetization networks in gaming rely on real-time data streaming to analyze player preferences and deliver targeted in-game advertisements and offers. This capability drives revenue growth through optimized ad spend and personalized upselling opportunities. By quickly adapting monetization tactics to player behaviors, gaming companies can sustain business growth and maximize profitability. The demand for robust infrastructure is further amplified by the rise of esports and interactive gaming, which require low latency and scalable architectures. Continuous server monitoring ensures 24/7 service health, rapid issue resolution, and seamless player xperiences.

These examples underscore the transformative role of data streaming in the gaming industry. Confluent enables real-time insights, personalized player experiences, and scalable operations to ensure that gaming companies remain agile and innovative in a competitive market. As the industry continues to evolve, continuous data processing will be pivotal in delivering engaging content, optimizing infrastructure, and achieving sustainable growth.

3.6 Government and Public Sector

Governments and public institutions are increasingly adopting digital public infrastructure to improve service delivery and enhance citizen experiences. The rise of smart cities and buildings is another critical trend, with IoT and data analytics enabling the creation of intelligent environments that improve quality of life and drive economic growth.

Environmental, Social, and Governance (ESG) considerations are also becoming more important, with public sector organizations focusing on sustainability and social responsibility.
The integration of data streaming platforms is revolutionizing the public sector by enabling real-time data processing and analytics.

Real-time capabilities in the public sector cannot be overstated, as it enables continuous monitoring, analysis, and sharing of data across institutions. The rise of digital public infrastructure and smart cities has further increased the demand for robust and scalable data streaming solutions.

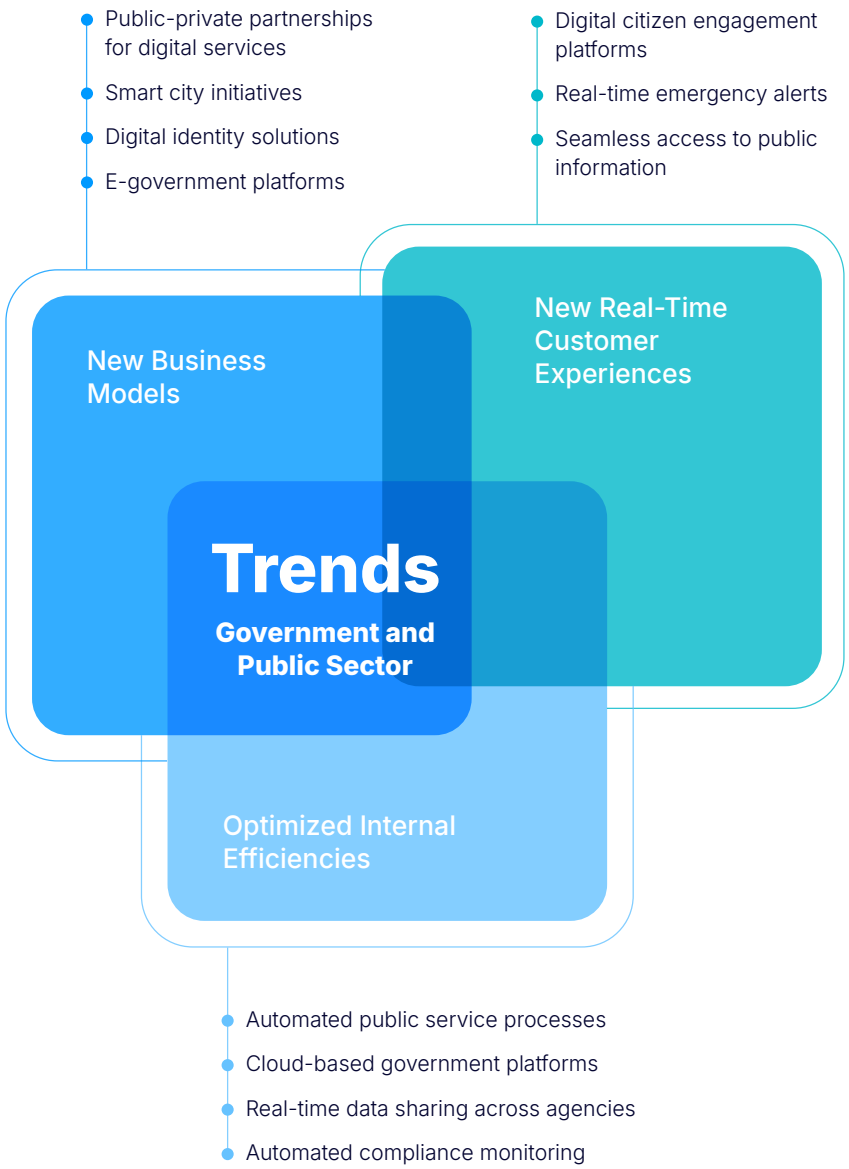

Scenario 1:
Citizen Services

Citizen services are present on various channels like online portals, mobile apps, and in-person interactions. Using real-time data from government, healthcare providers, and social services, organizations can create a single view of citizens, providing personalized and efficient services through all phases of life: work, family, health, retirement, and social security. Data streaming allows for a consistency that helps the public sector to improve inter-agency data sharing and enhance citizen experiences.

Business Value

- **Comprehensive citizen view:** Real-time data integration from multiple sources provides a single, comprehensive view of each citizen, enabling personalized services.
- **Improved inter-agency collaboration:** Enhanced data sharing between government agencies facilitates better coordination and collaboration, leading to more cohesive and effective public services.
- **Enhanced citizen experience:** Personalized and efficient service delivery improves the overall citizen experience, fostering trust and satisfaction with public sector organizations.
- **Streamlined service delivery:** Real-time data processing enables quicker response times and more efficient service delivery, reducing wait times and improving accessibility for citizens.
- **Data-driven decision-making:** Access to comprehensive citizen data supports informed decision-making, allowing public sector organizations to tailor services to meet the specific needs of their communities.
- **Increased operational efficiency:** By modernizing IT infrastructure and leveraging data streaming, public sector organizations can optimize their operations and reduce redundancies.
- **Proactive citizen support:** Real-time insights into citizen needs and circumstances allow for proactive support and intervention, improving outcomes in areas such as health, social security, and family services.
- **Cost savings:** Streamlined processes and improved data sharing reduce administrative costs and resource waste, leading to significant cost savings for public sector organizations.

Scenario 2:
Smart City

Smart cities use IoT and data analytics to create intelligent environments that improve quality of life and drive economic growth. Modern IT infrastructures and applications allow the use of real-time data from sources like traffic sensors, cameras, and connected vehicles to reduce traffic and road accidents through intelligent transportation systems. Data streaming enables smart cities to make data-driven decisions and enhance public safety.

Business Value

- **Improved traffic management:** Real-time data processing enables cities to monitor and manage traffic flow efficiently, reducing congestion and improving overall traffic conditions.
- **Enhanced public safety:** Intelligent transportation systems that leverage real-time data can reduce road fatalities and accidents by providing timely alerts and updates to drivers.
- **Data-driven decision-making:** Access to real-time traffic data supports informed decision-making, allowing city planners and officials to implement effective traffic management strategies.
- **Reduced environmental impact:** By optimizing traffic flow and reducing congestion, smart city initiatives can lower vehicle emissions, contributing to a cleaner and more sustainable urban environment.
- **Economic growth:** Improved transportation systems enhance the overall quality of life, attracting businesses and residents, and driving economic growth in the city.
- **Proactive infrastructure management:** Real-time insights into traffic patterns and infrastructure usage allow for proactive maintenance and management, extending the lifespan of city assets.
- **Innovation and modernization:** Implementing smart city technologies positions public sector organizations as leaders in urban innovation and modernization.

Scenario 3:
National Security

National security involves the protection of a country's citizens, infrastructure, and interests from various threats, including cyberattacks, terrorism, and natural disasters. Collection disciplines like IoT sensors, signals, and human intelligence, can continuously collect, analyze, and share data from vehicles, devices, apps, and infrastructure in real time. Leveraging these capabilities through data streaming allows government organizations to provide comprehensive, accurate, and current insights into the operating environment and enhance national security.

Business Value

- **Enhanced threat detection and response:** Real-time data processing enables the continuous monitoring and analysis of potential threats, allowing for rapid detection and response to cyberattacks, terrorism, and other security risks.
- **Comprehensive situational awareness:** By collecting and analyzing data from various sources, public sector organizations can gain a fuller picture of the operating environment, improving situational awareness and decision-making.
- **Proactive risk management:** Continuous data collection and analysis enable proactive risk management, reducing the likelihood of security breaches and enhancing overall resilience.
- **Inter-agency collaboration:** Enhanced data sharing between government agencies and departments facilitates better coordination and collaboration, leading to more cohesive and effective national security efforts.
- **Innovation and modernization:** Leveraging advanced technologies and real-time data analytics positions public sector organizations as leaders in security innovation and modernization.

Conclusion

Transformative trends, including smart city programs, e-government platforms, and public-private digital service partnerships, are driving the digital evolution of the public sector. IoT, data analytics, and real-time data streaming are reshaping how governments deliver services, interact with citizens, and manage operations. Real-time citizen engagement platforms, seamless access to public information, and automated compliance monitoring are enhancing efficiency and citizen satisfaction.

Smart city initiatives illustrate the transformative potential of real-time data. By integrating IoT devices, cities monitor traffic, optimize infrastructure, and improve public safety. Real-time traffic data reduces congestion, lowers emissions, and enhances urban quality of life, driving economic growth and sustainability. For example, intelligent transportation systems powered by data streaming provide timely alerts to reduce road accidents and fatalities. National security applications highlight the power of real-time data streaming in enhancing protection efforts. Governments gain comprehensive situational awareness and quicker threat responses through the analysis of data from sensors, signals, and human intelligence. Real-time analytics enable proactive risk management and seamless inter-agency collaboration to ensure resilience against cyberattacks, terrorism, and natural disasters.

Citizen services have been revolutionized with personalized, real-time interactions. Data streaming creates a unified view of citizens by integrating information from healthcare, social services, and government databases. This enables efficient, tailored service delivery across life phases, from healthcare to retirement. Enhanced inter-agency data sharing reduces redundancies and fosters collaboration, resulting in faster response times and significant cost savings.

These advancements position data streaming as a cornerstone of modern public sector operations. Continuous insights enable proactive governance, streamlined service delivery, and data-driven decision-making. By leveraging technologies like IoT and AI, governments can enhance citizen satisfaction, optimize resources, and drive sustainable development.

As public institutions evolve, integrating data streaming platforms will be critical for success. From improving urban mobility to safeguarding national security, real-time data processing empowers governments to respond dynamically to citizen needs, modernize infrastructure, and achieve greater transparency and efficiency in service delivery. These innovations underscore the public sector's ability to lead in a data-driven world while fostering trust and economic growth.

3.7 Energy and Utilities

One of the most significant trends in the energy and utilities sector is the push toward decarbonization and the integration of renewable energy sources. As countries and companies strive to meet ambitious climate goals, the adoption of solar, wind, and other renewable energy sources is accelerating. This shift necessitates the modernization of the grid to handle the intermittent nature of renewable energy and ensure reliable power delivery.

Grid modernization and the development of smart grids are critical components of this transformation. Smart grids leverage IoT, data analytics, and advanced communication technologies to provide real-time monitoring and control of the grid. This capability enhances grid reliability, optimizes energy usage, and enables the integration of distributed energy resources.

The digital customer experience is also becoming increasingly important in the energy and utilities sector. Customers now expect seamless, personalized interactions with their energy providers, similar to their experiences with other digital services. Data streaming platforms are playing a crucial role in enabling these trends by providing real-time data processing and analytics capabilities.

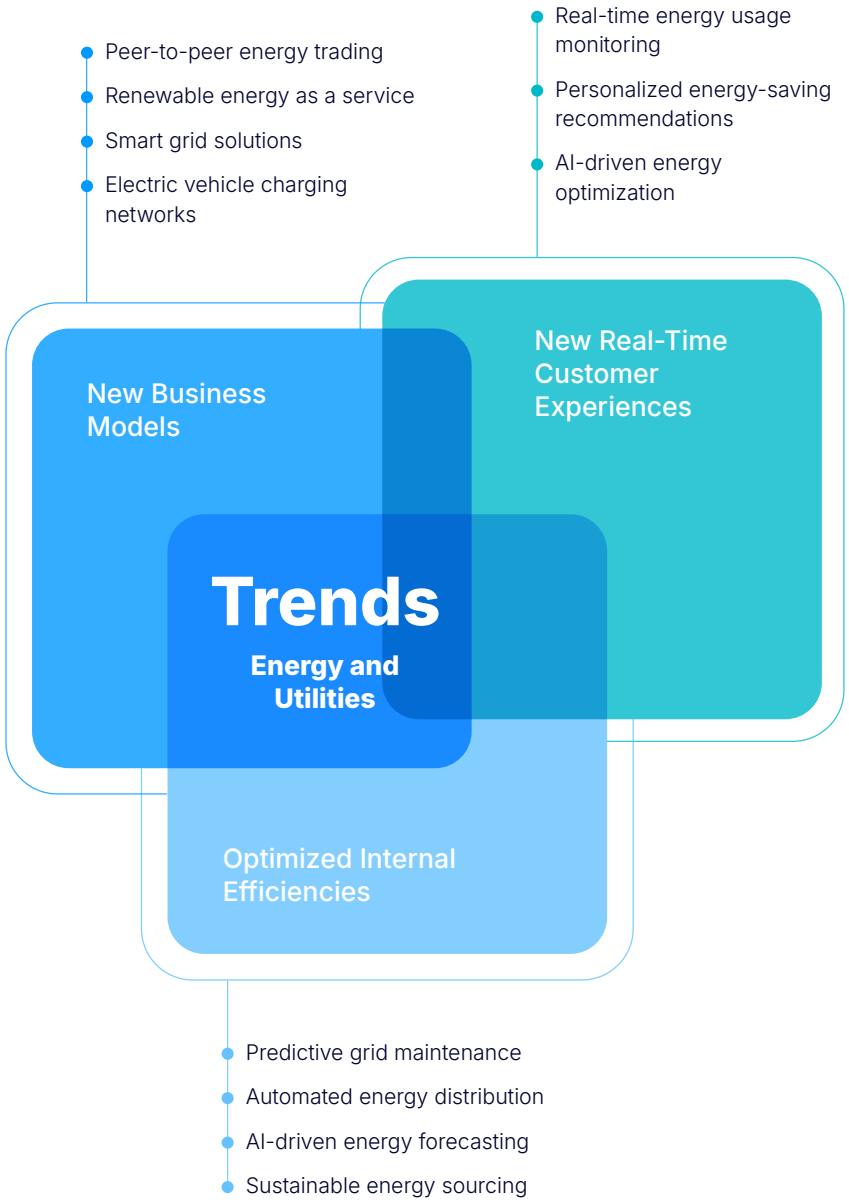

Scenario 1:
Smart Grid Monitoring

Smart grid monitoring ensures optimal performance and reliability by monitoring and managing the electrical grid, Data pipelines and IoT offer real-time insights into the grid to quickly identify and resolve issues. Data streaming allows working with real-time data from various sources like weather conditions, solar panel performance, and energy usage in order to quickly react to customer situations such as weather changes, tree shade, or component failures that impact power generation. This capability helps energy companies to optimize energy usage, improve grid reliability, and enhance customer satisfaction.

Business Value

- **Optimized energy usage:** Real-time data processing allows energy companies to monitor and manage energy consumption efficiently, optimizing energy usage and reducing waste.
- **Improved grid reliability:** Continuous monitoring of the electrical grid enables quick identification and resolution of issues, enhancing grid reliability and minimizing downtime.
- **Enhanced customer satisfaction:** By providing reliable energy services and timely responses to customer situations, energy companies can improve customer satisfaction and foster loyalty.
- **Comprehensive grid visibility:** Access to real-time data from various sources provides a comprehensive view of the grid, supporting informed decision-making and strategic planning.
- **Proactive maintenance:** Real-time insights into grid performance allow for proactive maintenance and management, reducing the likelihood of component failures and extending the lifespan of grid assets.
- **Sustainable energy management:** By optimizing energy usage and integrating renewable energy sources, smart grid monitoring supports sustainable energy management and environmental responsibility.
- **Competitive advantage:** Energy companies that leverage real-time data for smart grid monitoring can differentiate themselves by offering reliable and efficient energy services, gaining a competitive edge.

Scenario 2:
Industrial Control Systems

Industrial control systems (ICS) monitor and control industrial processes and equipment to ensure optimal performance and safety. Modern IT and IoT systems enable open, flexible, and elastic Supervisory Control and Data Acquisition (SCADA) systems that stream data in real time from sources like legacy OT gateways, ERP systems, and monitoring applications. These capabilities help energy companies massively to ensure business continuity and enhance operational efficiency.

Business Value

- **Business continuity:** Continuous data processing and monitoring capabilities guarantee business continuity by quickly identifying and resolving issues before they escalate.
- **Improved safety and compliance:** Real-time insights into industrial processes enhance safety measures and ensure compliance with industry regulations and standards.
- **Seamless integration:** Modernized SCADA systems that connect to legacy OT gateways, ERP systems, and modern cloud-native applications, providing seamless integration and data flow across the organization.
- **Proactive maintenance:** Access to real-time data allows for proactive maintenance and management of industrial equipment, reducing the likelihood of failures and extending asset lifespan.
- **Scalability and flexibility:** Open and flexible SCADA systems provide the scalability and flexibility needed to adapt to changing industrial demands and technological advancements.
- **Cost savings:** Improved operational efficiency and reduced downtime lead to significant cost savings, maximizing the return on investment in industrial control systems.
- **Innovation and modernization:** Leveraging IT modernization and IoT technologies positions energy companies as leaders in industrial innovation and modernization.

Scenario 3:
Energy Trading

Energy trading involves the buying and selling of energy commodities in wholesale markets. By leveraging data pipelines and IoT, energy companies can implement algo trading systems, energy trading and risk management (ETRM) systems, and external trading markets that are connected and data is correlated. This real-time data processing capability helps energy companies to process data faster, even with increased trading volumes and connected measurement devices.

By leveraging real-time data from various sources, including algo trading systems, ETRM systems, and external trading markets, energy companies can process and analyze trading data in real time. This real-time data processing capability allows organizations to process data faster and make data-driven trading decisions.

Business Value

- **Faster data processing:**
  Real-time data processing enables energy companies to handle increased trading volumes efficiently, ensuring timely and accurate data analysis.
- **Data-driven trading decisions:**
  Access to real-time data from various sources supports informed and strategic trading decisions, optimizing trading outcomes and profitability.
- **Enhanced risk management:**
  Integration of energy trading and risk management (ETRM) systems allows for continuous monitoring and assessment of market risks, improving risk mitigation strategies.
- **Increased trading efficiency:**
  Automated algo trading systems streamline trading operations, reducing manual intervention and increasing overall efficiency.
- **Improved market responsiveness:**
  Real-time insights into market conditions and trading data enable quick adaptation to market changes, enhancing competitiveness.
- **Seamless integration:**
  Connected data pipelines and IoT devices ensure seamless integration and data flow between trading systems and external markets.

Scenario 4:
Decarbonization

Decarbonization is a critical goal for the energy and utilities sector, driving the adoption of cleaner energy sources and more efficient energy management practices. This use case involves the deployment of IoT devices and data pipelines to collect and analyze data on energy usage, emissions, and environmental impact. By leveraging a data streaming platform, utilities can process and analyze this data in real time, enabling more efficient energy usage and improved ESG reporting.

Decarbonization involves the integration of IoT devices and data pipelines to collect and analyze data on energy usage, emissions, and environmental impact. This data is then processed and analyzed in real-time using data streaming. The insights gained from this analysis enable utilities to optimize their energy usage, reduce emissions, and improve their ESG reporting.

Business Value

- **Optimized energy usage:**
  Real-time data analysis allows utilities to monitor and manage energy consumption efficiently, reducing waste and optimizing energy usage.
- **Reduced emissions:**
  Continuous monitoring of emissions data enables utilities to identify and implement strategies to reduce carbon emissions, contributing to environmental sustainability.
- **Improved ESG reporting:**
  Access to comprehensive and real-time data supports accurate and transparent ESG (Environmental, Social, and Governance) reporting, enhancing corporate responsibility and reputation.
- **Enhanced regulatory compliance:**
  Real-time insights into energy usage and emissions ensure compliance with environmental regulations and standards, reducing the risk of penalties and fines.
- **Public trust and confidence:**
  By demonstrating a commitment to sustainability and environmental responsibility, utilities can build public trust and confidence in their operations.

Real-Life Examples

- AMPEERS ENERGY
- aedifion
- Uniper

Real-Life Example AMPEERS ENERGY

Use Case(s): Hybrid and Multi-Cloud | IoT

Company Overview

AMPEERS ENERGY is a leading provider of energy management solutions, specializing in the optimization of energy usage and the integration of renewable energy sources. The company offers a range of innovative solutions designed to help organizations achieve their decarbonization goals, improve energy efficiency, and enhance ESG reporting. AMPEERS ENERGY solutions leverage advanced IoT technologies and data streaming platforms to provide real-time insights and optimization capabilities.

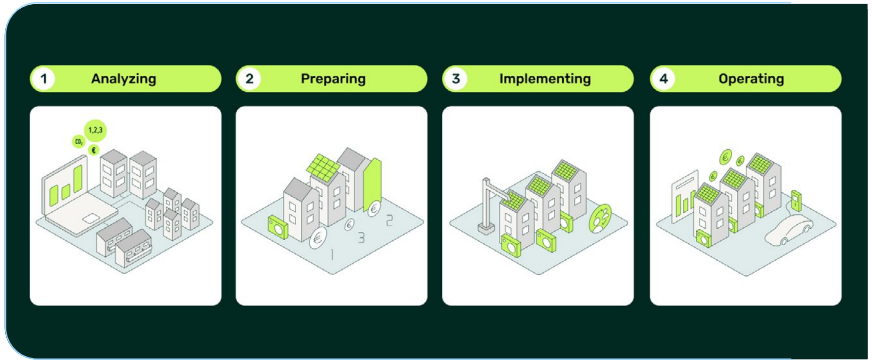

Implementation

AMPEERS ENERGY implemented a comprehensive decarbonization solution for a large real estate portfolio, focusing on district management and local supply optimization. The solution involved the deployment of IoT sensors and devices across multiple buildings to collect real-time data on energy usage, temperature, humidity, and other relevant parameters. This data was then streamed to a central platform powered by Confluent with Apache Kafka, where it was processed and analyzed in real time.

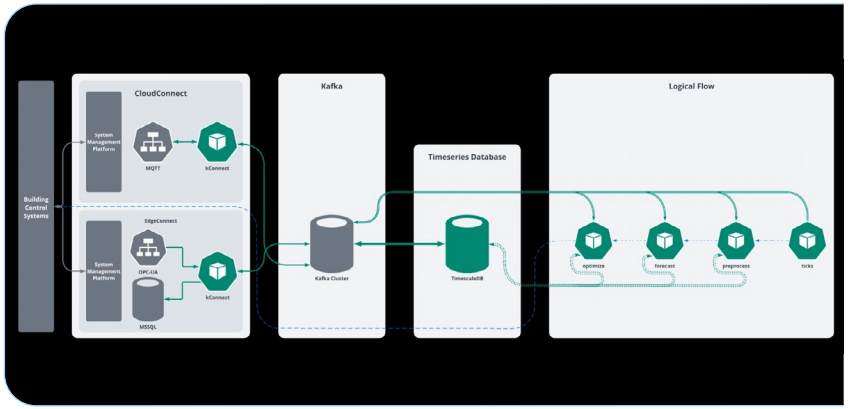

Architecture Overview of AMPEERS ENERGY's Kafka Implementation

The district management component of the solution involved the use of advanced forecasting and optimization algorithms to predict energy usage patterns and identify opportunities for optimization. By analyzing historical data and real-time inputs, the system could generate accurate forecasts of energy demand and implement strategies to reduce energy usage during peak periods. This included adjusting heating and cooling systems, optimizing lighting schedules, and implementing demand response measures.

The local supply management component focused on accounting for local energy usage and ensuring that energy was sourced from renewable or low-carbon sources. This involved integrating data from energy meters, weather forecasts, and renewable energy generation systems to provide a comprehensive view of energy supply and demand. By leveraging data streaming platforms, AMPEERS ENERGY was able to process and analyze this data in real time, enabling more informed decision-making and better management of energy resources.

"The biggest benefits for us to be in Confluent Cloud are that the management goes into someone that founded Kafka and then operates it on a well-known level. This takes responsibility away from us and we can simply say, 'This works'!"

Lucas Recknagel, *CTO, AMPEERS ENERGY*

Business Value

- **More efficient energy usage and reduced carbon emissions:** The implementation of the decarbonization solution allowed AMPEERS ENERGY to monitor and manage energy usage in real time, enabling more efficient energy consumption across multiple buildings. The district management component optimized energy usage, leading to substantial cost savings and reduced environmental impact. By leveraging advanced data analytics, AMPEERS ENERGY could identify inefficiencies and make data-driven decisions to further reduce energy consumption and emissions.

- **Improved renewable energy sourcing and support for decarbonization goals:** The technology ensures that energy is sourced from renewable or low-carbon sources, aligning with AMPEERS ENERGY and their clients' decarbonization objectives. This not only reduces reliance on traditional energy sources but also helpes AMPEERS ENERGY clients meet sustainability goals and comply with increasing regulatory demands for renewable energy usage.

- **Local supply management:** This feature (in German: "Mieterstrom") is a possibility to refinance investments via a local offered electricity tariff. The tariff chooses local produced electricity from solar panels over grid electricity. The solution offers landlords and tenants the economic benefits of quick and easy solar power generation.

- **Enhanced EMS reporting and regulatory compliance:** By using real-time data collection and analysis, AMPEERS ENERGY provided clients with accurate and comprehensive EMS reporting. The solution enabled precise tracking of energy usage and emissions, ensuring compliance with regulatory requirements and demonstrating clients' commitment to sustainability. The use of data streaming platforms allowed for the real-time processing of large data volumes, driving continuous improvement in energy management and long-term decarbonization efforts.

Real-Life Example aedifion

Use Case(s): Observability | IoT | ESG

Company Overview

aedifion is a leading provider of cloud-based energy management solutions for optimizing building portfolios. The company's platform leverages IoT sensors, data streaming technologies, and advanced analytics to monitor and optimize building performance in real time. aedifion's solutions are designed to help property managers reduce energy consumption, lower operational costs, and achieve sustainability goals.

Implementation

aedifion implemented a comprehensive energy management solution for a large real estate portfolio, integrating with the existing building management systems (BMS) to provide real-time data and insights. The solution involved deploying IoT sensors throughout the buildings to monitor various parameters, including temperature, humidity, occupancy, and energy usage. These sensors generated vast amounts of data that were streamed in real time to aedifion's cloud platform using Confluent.

The data streaming platform facilitated the real-time processing and analysis of the sensor data, enabling aedifion to monitor building performance continuously. Advanced analytics and machine learning algorithms were applied to identify inefficiencies, detect anomalies, and recommend optimization measures. The platform also provided property managers with real-time dashboards and alerts, allowing them to take immediate action to address issues and optimize energy usage.

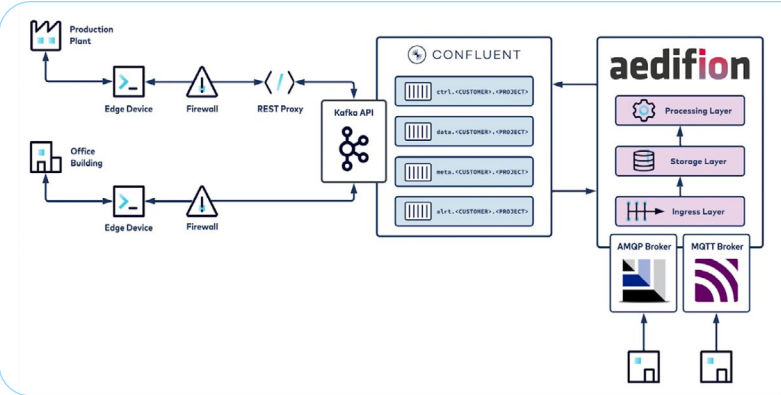

Confluent-based data pipeline at aedifion

Business Value

- **Energy savings, reduced operational costs, lower carbon footprint:** The decarbonization initiative provided significant business value for aedifion's clients, like for example 17.1 % CO_2 savings. By leveraging real-time data streaming and advanced analytics, the energy management solution enabled property managers to optimize building operations and enhance energy efficiency.
- **Preventing energy waste and improving overall sustainability:** The real-time visibility into building performance allowed property managers to identify and address inefficiencies promptly, preventing energy waste and improving overall sustainability. The solution also supported compliance with environmental regulations and sustainability certifications, enhancing the reputation and marketability of the real estate portfolio.
- **Improved comfort and satisfaction of building occupants:** Furthermore, the integration of IoT sensors and data streaming technologies enabled continuous monitoring and optimization, ensuring that buildings operated at peak efficiency at all times. This proactive approach to energy management not only contributed to decarbonization efforts but also improved the comfort and satisfaction of building occupants.

Significant reductions achieved in the Kaiser Hof Cologne

Real-Life Example Uniper

Use Case(s): Event-Driven Microservices | Data Sharing

Company Overview

Düsseldorf-based Uniper is a European energy company with global reach and activities in more than 40 countries. With approximately 7,400 employees, the company makes an important contribution to security of supply in Europe, particularly in its core markets of Germany, the UK, Sweden and the Netherlands. Uniper's operations encompass power generation in Europe, global energy trading, and a broad gas portfolio. Uniper procures gas—including liquefied natural gas (LNG)—and other energy sources on global markets. The company owns and operates gas storage facilities with a total capacity of more than 7 billion cubic meters.

Uniper intends to be completely carbon-neutral by 2040. Uniper aims for its installed power generating capacity to be more than 80% zero-carbon by the early 2030s. To achieve this, the company is transforming its power plants and facilities and investing in flexible, dispatchable power generating units. Uniper is already one of Europe's largest operators of hydropower plants and is helping further expand solar and wind power, which are essential for a more sustainable and secure future. The company is progressively expanding its gas portfolio to include green gases like hydrogen and biomethane and aims to convert to these gases over the long term.

Uniper is a reliable partner for communities, municipal utilities, and industrial enterprises for planning and implementing innovative, lower-carbon solutions on their decarbonization journey. Uniper is a hydrogen pioneer, is active worldwide along the entire hydrogen value chain, and is conducting projects to make hydrogen a mainstay of the energy supply.

Implementation

In the fast-paced and complex world of energy trading, Uniper recognized the need to leverage advanced technologies to optimize its trading operations and maintain a competitive edge. The wholesale trading market involves a myriad of interactions between energy producers, utilities, and traders, all working to balance energy supply and demand. To navigate this intricate ecosystem, Uniper implemented data streaming technologies to enable real-time data processing and analytics, thereby enhancing its trading capabilities.

Uniper's energy trading operations are underpinned by Confluent's data streaming platform. This platform facilitates the continuous streaming of data from a variety of sources, including market prices, weather forecasts, and grid conditions. By processing this data in real time, Uniper supports its algo trading systems, energy trading and risk management (ETRM) systems, and external trading markets.

The data streaming platform enables seamless integration and correlation of data, providing traders with real-time insights and actionable information. This integration is crucial for creating a seamless trading environment where algo trading systems, ETRM systems, and scheduling/dispatch/invoice systems are interconnected.

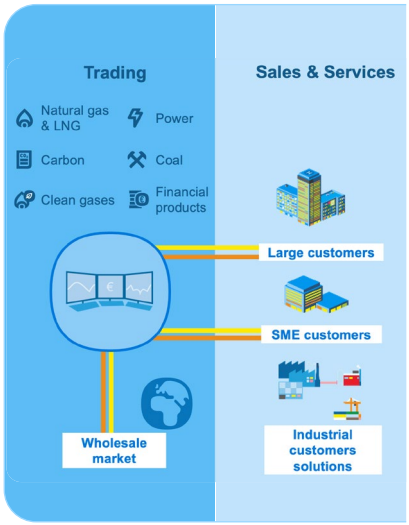

Uniper's Sales & Trading focus market areas

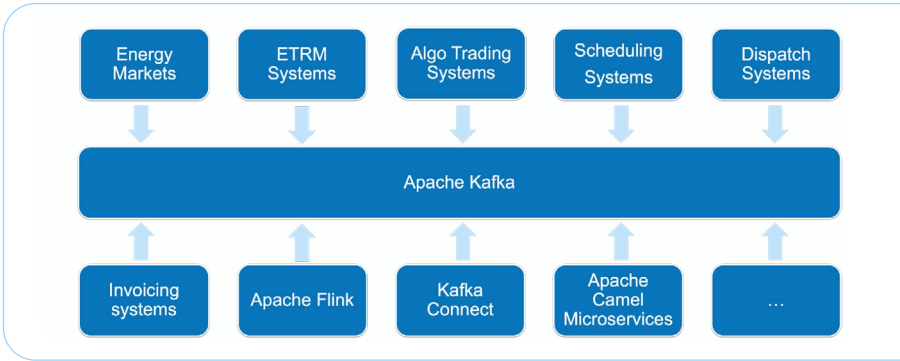

Decoupled producers and consumers at Uniper

Uniper applies advanced analytics and machine learning algorithms to the streamed data to optimize trading strategies, manage risks, and enhance decision-making. Real-time dashboards and alerts are employed to allow traders to continuously monitor market conditions and respond swiftly to changes. This capability ensures that Uniper's trading operations are not only profitable but also efficient.

Business Value

- **Faster data processing:** The ability to process data in real time enables Uniper to respond quickly to market changes, optimizing trading strategies and capturing opportunities as they arise.
- **Increased trading volumes:** The data streaming platform supports increased trading volumes and the integration of several trading systems. This allows Uniper to expand its trading activities.
- **Enhanced decision-making:** Real-time insights and advanced analytics improve trading decisions and risk management, providing Uniper with a competitive advantage in the energy market.
- **Operational efficiency:** The seamless integration of trading systems and data sources enhances operational efficiency to reduce the complexity and cost of managing trading operations.
- **Scalability and flexibility:** The cloud-based data streaming platform provides the scalability and flexibility needed to handle large volumes of data and support Uniper's business growth. This scalability ensures that Uniper can adapt to changing market conditions and continue to innovate in its trading operations.

"We saw the demand in our business rising, trading volumes have been increasing significantly, and this put more pressure on our systems. More messages were coming in daily and we required faster processing of data. We did not want to miss any opportunity on the market just because our IT processes were not fast enough. The decision to go with data streaming based on Apache Kafka and Apache Flink was both business and architecture driven."

Alexander Esseling, *Platform & Security Architect, Uniper*

Conclusion

The energy and utilities sector is undergoing a significant transformation driven by decarbonization, renewable energy integration, and advancements in digital technologies. New business models such as peer-to-peer energy trading, renewable energy as a service, and electric vehicle charging networks are reshaping the industry, while real-time energy usage monitoring and AI-driven optimization enhance customer experiences. redictive grid maintenance, automated energy distribution, and AI-driven forecasting improve operational efficiency and sustainability.

AMPEERS ENERGY highlights the impact of continuous data streaming in decarbonization efforts. By deploying IoT sensors across real estate portfolios, AMPEERS processes real-time data to optimize energy usage, reduce emissions, and enhance ESG reporting. This approach not only achieves significant energy savings but also aligns with regulatory and sustainability goals.
aedifion utilizes data streaming to enhance property efficiency, combining IoT sensors with advanced analytics to monitor energy use and minimize waste. Its solutions elevate tenant well-being and sustainability while ensuring adherence to environmental standards.

Uniper exemplifies the power of data streaming in energy trading. By processing real-time data from energy markets, weather forecasts, and grid conditions, Uniper enables algo trading and risk management. This capability allows Uniper to increase trading volumes, improve decision-making, and maintain operational efficiency. Real-time dashboards provide actionable insights to ensure swift responses to market changes and optimized trading strategies.

Smart grid monitoring further showcases the transformative potential of data streaming. By leveraging IoT devices and real-time analytics, energy providers can enhance grid reliability, optimize energy distribution, and proactively manage maintenance. These capabilities reduce downtime, improve customer satisfaction, and support the integration of renewable energy sources.

These advancements illustrate the critical role of data streaming in modern energy and utility operations. Continuous data processing enables companies to optimize grid performance, deliver personalized customer experiences, and meet ambitious decarbonization targets. As exemplified by AMPEERS ENERGY, aedifion, and Uniper, data streaming empowers the energy sector to innovate, improve efficiency, and drive sustainable growth in an evolving market.

3.8 Transportation and Travel

The transportation and travel industry is experiencing a wave of digitalization and innovation, driven by advancements in technology and changing consumer expectations.

One of the most significant trends is the adoption of robotics and autonomous vehicles, which are revolutionizing the way goods are transported and delivered. These technologies enhance operational efficiency, reduce costs, and improve safety. Biometric and digital identity solutions are also becoming increasingly important in the travel industry. These technologies enable seamless and secure travel experiences, from check-in to boarding, by leveraging biometric data for identity verification.

Improving the B2B and B2C customer experience is another critical trend in the logistics and travel industry. Customers now expect personalized, seamless interactions with logistics and travel providers, similar to their experiences with other digital services. Data streaming plays a crucial role in enabling these trends by providing real-time data processing and analytics capabilities.

The rise of robotics, autonomous vehicles, and digital identity solutions has further increased the demand for robust and scalable data streaming solutions. As the logistics and travel industry continues to evolve, the adoption of data streaming technologies will be crucial in maintaining competitiveness and driving innovation.

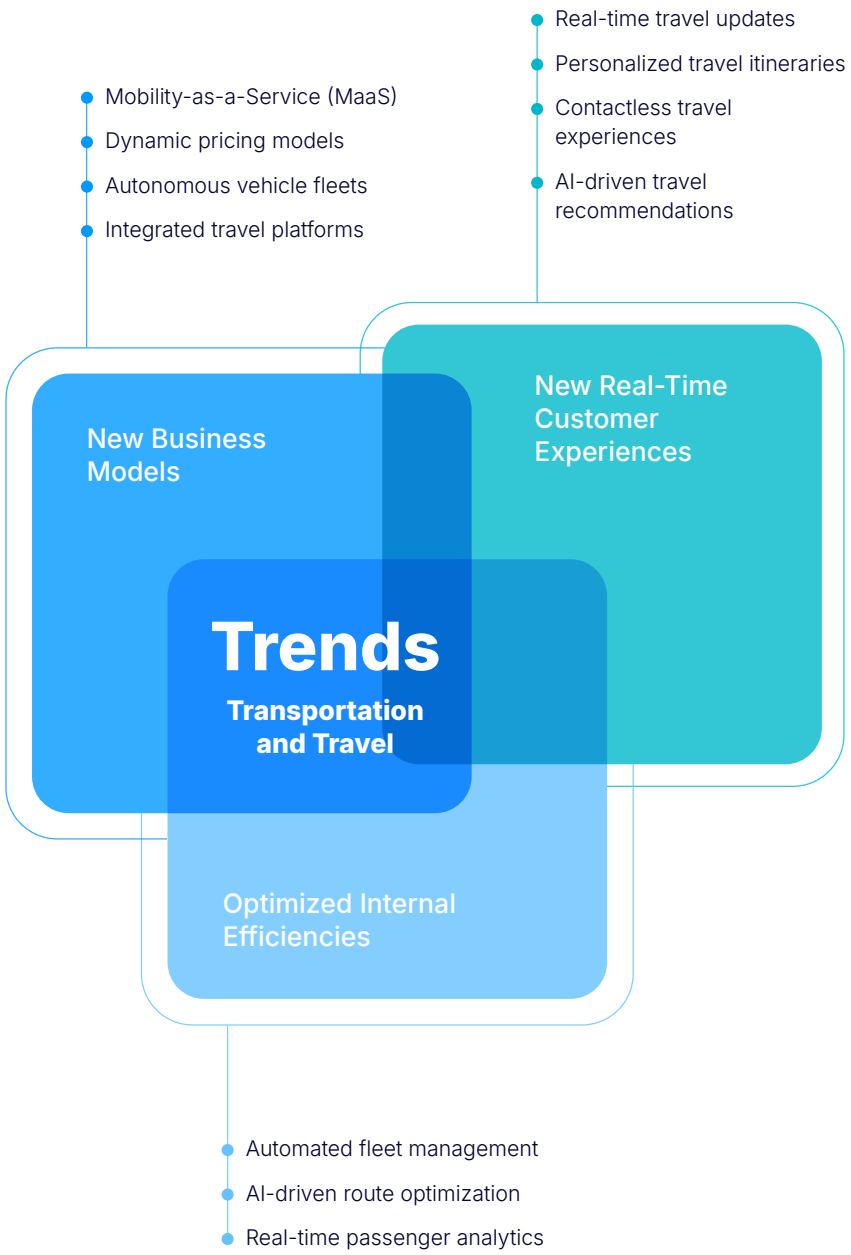

Scenario 1:
Transport Management

Transport management involves the planning, execution, and monitoring of the movement of goods and people. IoT and ESG (Environmental, Social, and Governance) principles can help logistics companies transform their transport management systems (TMS) into transport visibility & management systems (TVMS). Data streaming can be used to monitor the movement of vehicles and goods in real time and provide real-time updates to customers, providing end-to-end visibility for logistics processes, optimizing operations, and reducing environmental impact.

Business Value

- **End-to-end visibility:** Real-time data processing provides comprehensive visibility into the logistics process, allowing companies to monitor the movement of goods and people effectively.
- **Optimized logistics operations:** Continuous monitoring and analysis of transportation data enable companies to optimize routes, reduce delays, and improve overall logistics efficiency.
- **Reduced environmental impact:** By integrating ESG principles and leveraging real-time environmental data, companies can implement strategies to minimize their carbon footprint and promote sustainability.
- **Enhanced customer satisfaction:** Real-time updates and transparency in logistics operations improve customer satisfaction by providing timely and accurate information on the status of shipments.
- **Data-driven decision-making:** Access to real-time data from IoT sensors, GPS tracking, and environmental sources supports informed decision-making, allowing companies to implement effective transport management strategies.

Scenario 2:
Ride-Hailing

Ride-hailing provides on-demand transportation services through mobile applications. Data streaming is implemented in these services to optimize operations, connecting real-time data from various sources, including driver availability, customer demand, and traffic conditions. These capabilities offer a comprehensive view of ride-hailing operations, letting ride-hailing companies implement dynamic pricing, real-time campaigns, and fraud detection systems.

Business Value

- **Optimized pricing strategies:** Real-time data processing enables dynamic pricing models that adjust fares based on current demand, driver availability, and traffic conditions, maximizing revenue and market competitiveness.
- **Enhanced customer satisfaction:** By providing timely and efficient ride-hailing services, companies can improve customer satisfaction and foster loyalty through personalized and responsive service offerings.
- **Maximized resource utilization:** Real-time insights into driver availability and customer demand allow for optimal allocation of resources, ensuring drivers are available where and when they are needed most.
- **Effective fraud detection:** Continuous monitoring and analysis of ride-hailing data support the implementation of robust fraud detection systems, protecting both the company and its customers from fraudulent activities.
- **Data-driven marketing campaigns:** Real-time data supports the execution of targeted marketing campaigns, enhancing customer engagement and driving demand during peak times.
- **Competitive advantage:** Companies that implement advanced ride-hailing systems can differentiate themselves by offering reliable and efficient services, gaining a competitive edge in the market.

Scenario 3:
Travel Booking

Travel booking involves the planning and booking of travel arrangements, including flights, hotels, and activities. Data pipelines and AI/ML allow travel companies to integrate chatbots and generative AI to support customer service functions such as rebooking, refunds, and itinerary changes. With data streaming, these systems can access real-time data from various sources, including customer interactions, booking data, and travel information. These capabilities make it possible for travel companies to provide automated self-service options, enhance customer satisfaction, and increase net promoter scores.

Business Value

- **Enhanced customer satisfaction:** By providing efficient and responsive customer service through chatbots and generative AI, travel companies can improve customer satisfaction and foster loyalty.
- **Automated self-service options:** Real-time data processing and AI integration enable automated self-service options, allowing customers to manage bookings, rebook, request refunds, and make itinerary changes independently.
- **Increased net promoter scores:** Streamlined and personalized customer interactions contribute to higher net promoter scores, reflecting improved customer experiences and brand advocacy.
- **Comprehensive operational visibility:** Access to real-time data from customer interactions, booking data, and travel information provides a comprehensive view of vacation booking operations, supporting informed decision-making.
- **Improved operational efficiency:** Automated processes and efficient resource management reduce operational costs and improve the overall efficiency of vacation booking services.
- **Data-driven personalization:** Real-time insights into customer preferences and behaviors enable personalized service offerings, enhancing customer engagement and satisfaction.

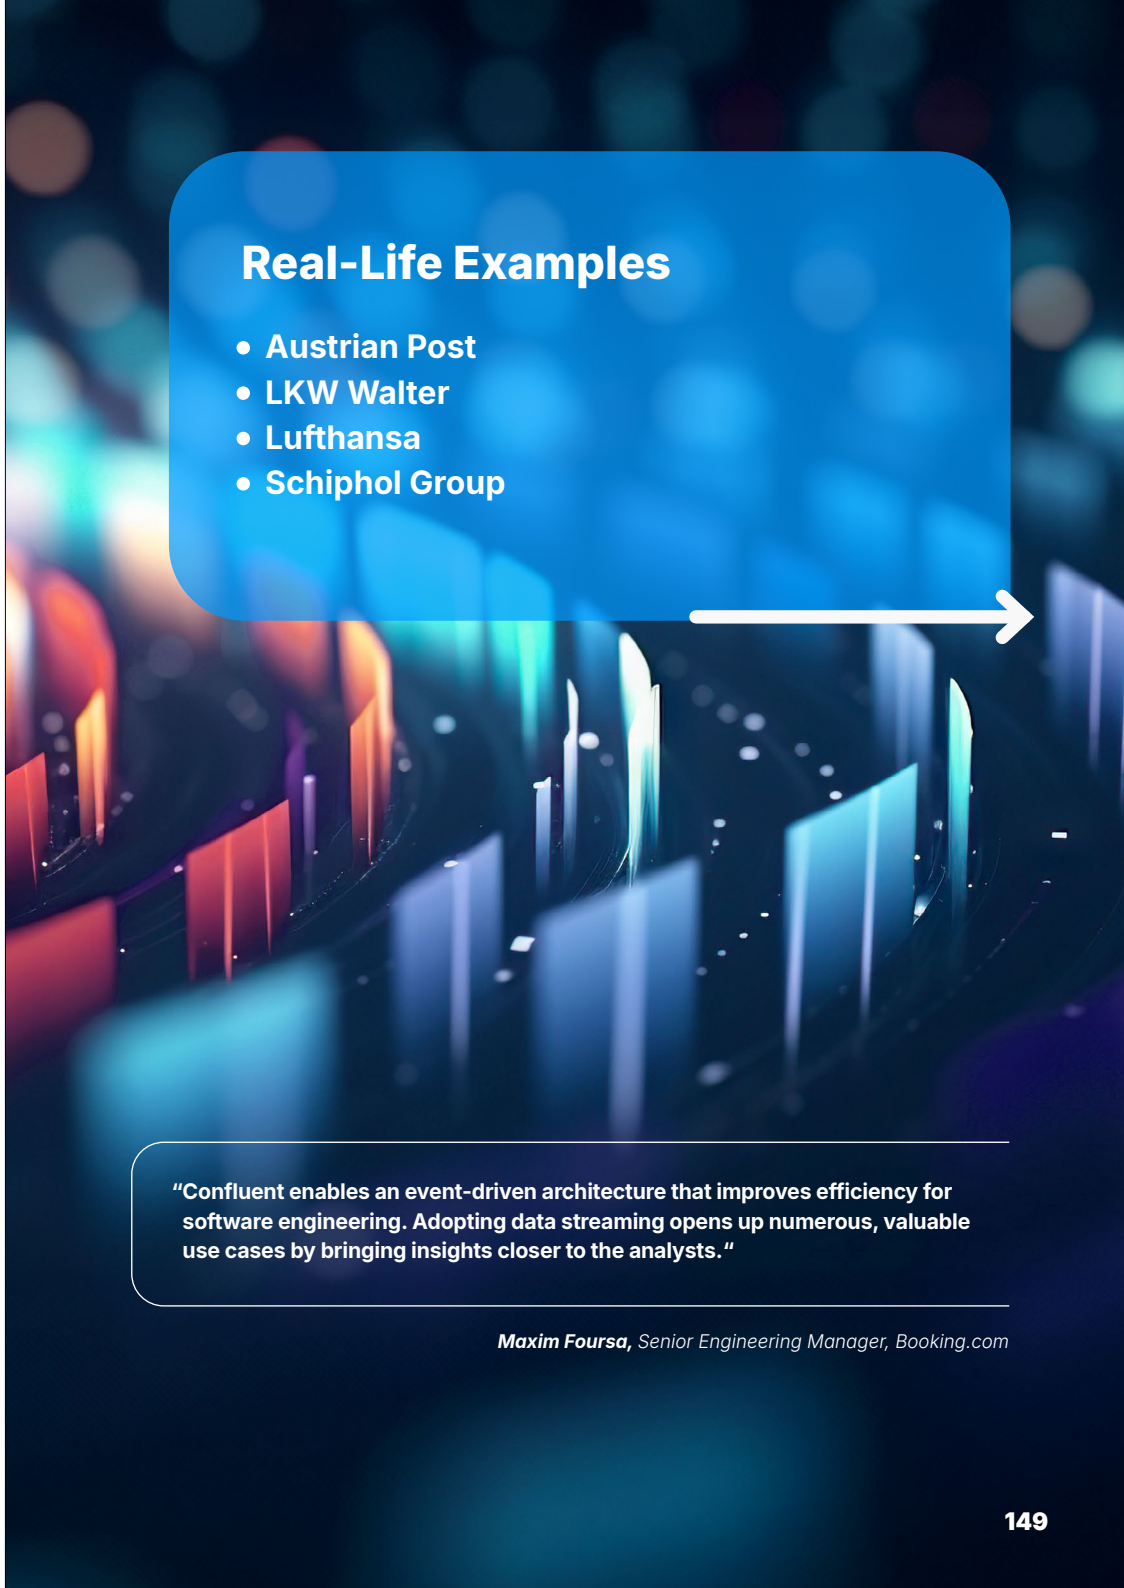

Real-Life Examples

- **Austrian Post**
- **LKW Walter**
- **Lufthansa**
- **Schiphol Group**

"Confluent enables an event-driven architecture that improves efficiency for software engineering. Adopting data streaming opens up numerous, valuable use cases by bringing insights closer to the analysts."

Maxim Foursa, *Senior Engineering Manager, Booking.com*

Real-Life Example Austrian Post

Use Case(s): Data Products | Observability | Event-Driven Microservices

Company Overview

Austrian Post is an international postal, logistics, and service provider which is central to Austria's economy. The company is firmly focused on the very highest quality and offers a comprehensive product and service portfolio to provide the best possible match for current customer needs. Austrian Post groups its operations into three divisions: Mail, Parcel & Logistics, and Retail & Bank. The company also has an international presence, operating in the markets of Germany, Southeast and Eastern Europe, Turkey, and Azerbaijan.

Implementation

The Austrian Post faced challenges with its existing shipment data management system ('track-and-trace'). The data was centralized in a large and conventional relational database system. It was difficult to maintain the system due to its monolithic structures. This, in turn, led to scaling limitations. Additionally, the fact that many other systems were dependent on the shipment data meant an ever-increasing load was placed on the core system.

There was no conceivable migration scenario beyond a big-bang replacement, which, for obvious reasons, was not an option. To address these issues, the Austrian Post implemented a comprehensive solution involving their Confluent-based data streaming platform, a newly built façade around their legacy core system, and a GraphQL-based item search to provide easy access to the shipment data.

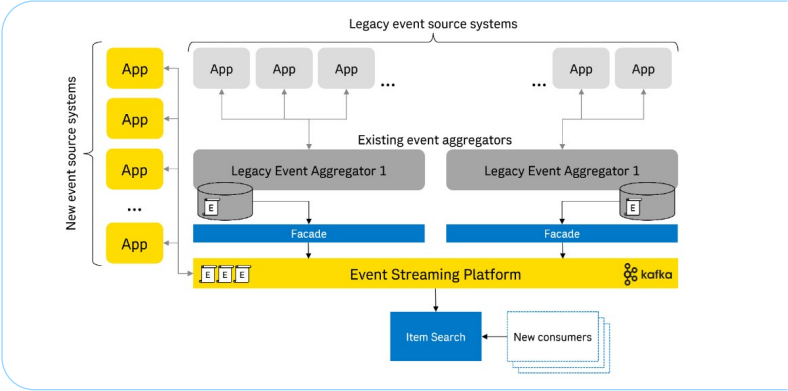

The facade wraps the legacy monolith(s)

The solution involves continuously pulling data from the legacy database, enhancing it with unique IDs, organizing it into domain-specific aggregates, and making it easily searchable and accessible through an API. This solution leverages Kafka-based technology, whereby a couple of dozen complex stream applications are employed to handle high-volume data, reduce monolith dependencies, allow access to important data without scaling issues, and provide a way to merge data for specific domains that are currently spread over several legacy applications. Additionally, it marks the first important step in a sweeping transformation strategy that will allow the Austrian Post to decompose the monolith(s) in the future.

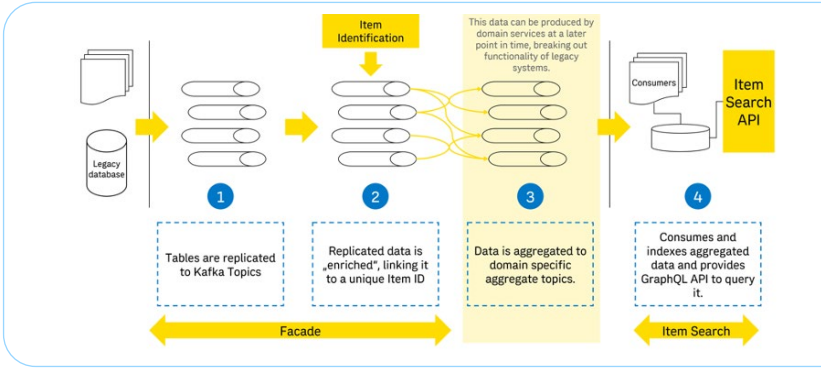

From legacy database to scalable item search

Business Value

- **Increased transparency:** The real-time availability of shipment data increases internal transparency and enables full item visibility with additional scans.
- **Increased customer satisfaction:** Rapid access to their own shipment data at scale enhances the customers' experience, resulting in higher satisfaction levels.
- **Reduced time to market:** The solution provides a shortcut to access data without waiting for all the individual systems to connect to the event streaming platform.
- **Improved cross-product visibility:** The ability to merge domain data from legacy applications increases cross-product visibility and simplifies access to data.

Real-Life Example LKW Walter

Use Case(s): Observability | Event-Driven Microservices | Data Sharing

Company Overview

LKW Walter, a leading European transport organization, implemented a comprehensive TMS leveraging Confluent as a data streaming platform. The solution involved deploying IoT sensors on trucks and throughout the supply chain to monitor their location and status in real time. These sensors generated vast amounts of data that were streamed in real time to LKW Walter's cloud platform using Confluent.

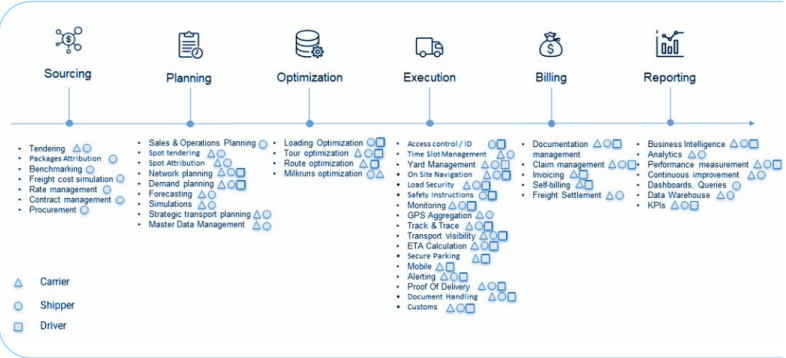

The end-to-end logistics process at LKW Walter that is now fully event-driven

Implementation

The data streaming platform facilitated the real-time processing and analysis of the sensor data, enabling LKW Walter to monitor and optimize truck logistics operations continuously. Advanced analytics and machine learning algorithms were applied to optimize routes, reduce fuel consumption, and ensure that deliveries were made on time and within budget. The platform also provided real-time dashboards and alerts, allowing logistics managers to take immediate action to address issues and optimize operations.

> **"An event-driven architecture lets us respond to events in real time, dramatically boosting the efficiency of our business processes. By automating decisions and notifications instead of relying on manual intervention, we keep customers updated on the status of their deliveries in real time. This level of agility and transparency gives us a distinct competitive edge."**

Leo Hintersteiner, *CIO, LKW Walter*

Predictive visibility as a service at LKW Walter

Business Value

- **Reduced transportation costs:** By optimizing routes and reducing fuel consumption, LKW Walter was able to significantly lower transportation costs.
- **Enhanced delivery reliability:** Real-time monitoring and optimization ensured that deliveries were made on time, enhancing reliability.
- **Improved operational efficiency:** Continuous monitoring and real-time data processing enabled efficient management of logistics operations.
- **Increased customer satisfaction:** Timely deliveries and improved reliability led to higher customer satisfaction.
- **Competitive advantage:** Efficient truck logistics operations provided LKW Walter with a competitive edge in the market.

As a concrete measurable example, LKW Walter receives double the stall fees compared to before, 2x more stall money thanks to real-time visibility as the company can now charge accurate penalties.

Real-Life Example Lufthansa

Use Case(s): IT Modernization | Data Sharing | AI

Company Overview

Lufthansa is a major German airline and one of the largest in Europe, renowned for its extensive network of domestic and international flights. As a member of the Star Alliance, Lufthansa offers a wide range of services, including passenger transportation and cargo logistics.

Implementation

Traditionally, airlines have used message queuing (MQ), enterprise service buses (ESB), and extract-transform-load (ETL) tools for system integration. However, these methods often lead to tightly coupled systems, limiting flexibility and scalability. Lufthansa has adopted data streaming as a cloud-native middleware, capitalizing on its real-time messaging, persistence layer, data integration, and stream processing capabilities. This approach decouples systems, manages backpressure, and enables historical event replay, enhancing both operational and analytical data integration.

Lufthansa's KUSCO (Kafka Unified Streaming Cloud Operations) project underscores its commitment to data streaming for mission-critical operations. By choosing Confluent over traditional messaging queues, Lufthansa achieved a more cost-effective and scalable infrastructure for real-time data processing. The KUSCO platform (now known as One Integration Platform) has streamlined project onboarding and provides real-time, high-quality data feeds across business units.

Using data streaming, Lufthansa drives real-time analytics and AI/machine learning applications. By consolidating and processing data from various sources, the airline enables real-time alerting and model scoring, delivering timely and accurate predictions for operational decision-making.

Use cases include for example:

1. **Anomaly detection:** Lufthansa uses stream processing to aggregate data and detect anomalies in real time, enabling proactive alerts and maintenance.
2. **Fleet management:** The airline leverages machine learning models embedded within the data streaming applications to optimize aircraft operation to ensure efficient fleet management with low latency and critical SLAs.

Machine Learning for Real-Time Operations Steering at Lufthansa with a high-level architecture

Business Value

The integration of data streaming technologies has delivered substantial business value to Lufthansa:

- **Operational efficiency:** By decoupling systems and enabling real-time data processing, Lufthansa has enhanced its operational efficiency, reducing costs and improving scalability.
- **Innovation and competitiveness:** Data streaming has empowered Lufthansa to innovate rapidly, offering personalized customer experiences and maintaining a competitive edge in the aviation industry.
- **Enhanced customer experience:** Real-time analytics and AI applications have enabled Lufthansa to provide superior customer service, from predictive maintenance to personalized travel experiences.
- **Scalability and flexibility:** The cloud-native architecture of the One Integration Platform allows Lufthansa to scale its operations efficiently, accommodating growing data volumes and diverse use cases.
- **Data-driven decision-making:** Access to real-time data feeds has facilitated data-driven decision-making across Lufthansa's business units, driving continuous improvement and strategic growth.

Real-Life Example Schiphol Group

Use Case(s): IT Modernization | Data Sharing | IoT | ESG

Company Overview

Schiphol Group, a prominent Dutch company, is renowned for its management and operation of airports in the Netherlands and internationally, with Amsterdam Airport Schiphol being its flagship. As one of Europe's busiest and most significant airports, Schiphol is a hub of activity, facilitating millions of passengers and tons of cargo annually. The Schiphol Group's operations extend beyond aviation services to encompass a wide range of non-aviation activities, making it a comprehensive airport management entity.

In its quest to become a leading autonomous airport by 2050, Schiphol Group has embarked on a digital transformation journey. This transformation is driven by the need to enhance operational efficiency, improve passenger experiences, and ensure seamless integration of various airport functions. The adoption of data streaming has been central to this transformation, enabling real-time data processing, data exchange and data analysis across multiple facets of airport operations.

Implementation

The digitalization of Schiphol Airport is a multifaceted endeavor that leverages data streaming to revolutionize airport operations and passenger experiences. At the core of this transformation is the implementation of data streaming, which facilitates the continuous flow and immediate processing of critical data from diverse sources such as flight information, passenger movements, and physical asset performance.

Event-driven is becoming more and more the basic principle for everyday data processing operations within Schiphol. It enables them to process incoming events in real time, providing transparency and context-specific decision-making capabilities not only within the organization but even with APIs to the outside world.

This architecture supports a wide range of use cases coming to life in the future, from predictive maintenance and personalized passenger services to enhanced security measures.

One of the key challenges addressed by Schiphol's digitization efforts is correlating trusted and qualitative data from different sources and departments. This will enable Schiphol to optimize operational excellence and realize the integration of Internet of Things (IoT) scenarios. Data sources can be events

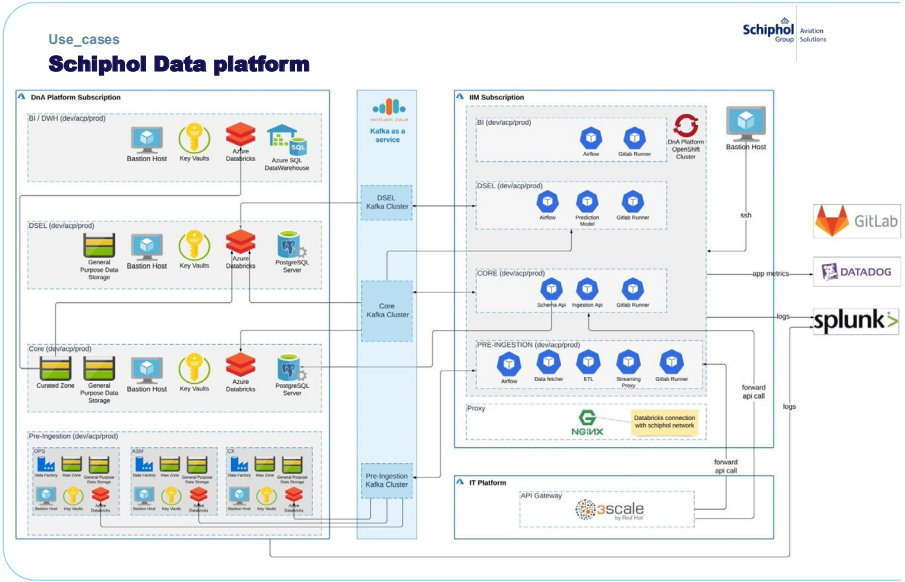

Schiphol Data Platform

from various interfaces, including gates, retail stores, and passenger mobile apps. By continuously correlating data in real time, Schiphol can implement predictive forecasting, planning, and maintenance, as well as develop advertisement and recommendation engines to improve customer experience and increase revenue.

The implementation of data streaming technologies has also facilitated significant improvements in operational efficiency. Automation with IoT sensors, paperless processes, and software innovations have enabled more cost-effective and reliable airport operations and cargo processes like flow management systems or platforms for skilled professionals to coordinate a smooth operation of the airport, while predictive maintenance powered by IoT sensors and data analytics ensures smoother operations by anticipating equipment failures before they occur.

Schiphol Group has modernized its legacy IT infrastructure by complementing traditional systems like message queues, ETL processes, and ESBs with a data streaming platform. This transition has enabled scalable and cloud-native data processing and seamless integration across diverse airport systems, significantly reducing scalability and data availability issues and enhancing operational efficiency.

Business Value

- **Enhanced operational efficiency:** Reduced end-to-end latency in data processing and automation of key processes have streamlined airport operations (e.g., electronic cargo documents, reduced turnaround times), minimizing disruptions and optimizing coordination.
- **Improved passenger experience:** Mobile apps and boarding, providing passengers with a seamless, efficient, and therefore satisfying journey.
- **Elevated customer service and communication:** Real-time updates enhancing passenger satisfaction and loyalty engagement.

Conclusion

The transportation and travel industry is rapidly evolving, shaped by trends like Mobility-as-a-Service (MaaS), autonomous vehicle fleets, and integrated travel platforms. Data streaming enables businesses to deliver innovative services, meeting customer expectations for real-time updates, personalized itineraries, and seamless, contactless experiences. Operational efficiency is elevated through AI-driven route optimization, automated fleet management, and real-time passenger analytics.

Austrian Post highlights the transformative potential of data streaming for shipment tracking. Real-time visibility and reduced latency streamline logistics processes, improving customer satisfaction and enabling faster service rollouts. Similarly, LKW Walter employs IoT sensors to monitor fleets in real-time, optimizing routes, lowering transportation costs, and ensuring reliable deliveries.

In aviation, Lufthansa integrates real-time data for predictive maintenance, anomaly detection, and fleet management. This ensures operational efficiency while delivering personalized services and improving customer engagement. Schiphol Group takes a holistic approach to airport digitalization with an event-driven architecture that optimizes operations, enhances passenger experiences, and supports predictive maintenance.

Real-time data streaming serves as a critical enabler for companies across transportation and travel to streamline processes, personalize services, and scale operations efficiently. By adopting event-driven architectures, businesses can integrate IoT, AI, and analytics to deliver seamless, customer-centric services while staying ahead of operational and cybersecurity challenges in a competitive landscape.

3.9 Insurance

When mentioning “insurance“ and “digitalization,“ “regulation“ quickly follows. This impacts the use of data and the adoption of certain trends at every stage. One of the most prominent trends is the adoption of telematics and IoT, which enable insurers to collect real-time data on customer behavior and asset conditions. This data is crucial for developing personalized insurance products and improving risk assessment. Digital customer engagement is another critical trend, as customers now expect seamless, personalized interactions with their insurers, similar to their experiences with other digital services.

Usage-based insurance (UBI) is gaining traction, allowing insurers to offer policies based on actual usage rather than traditional risk factors. This approach is particularly relevant in the automotive sector, where telematics data can be used to monitor driving behavior and adjust premiums accordingly. Embedded insurance, which integrates insurance products into the purchase process of other goods and services, is also becoming increasingly popular. This trend simplifies the insurance buying process and enhances customer convenience.

Data streaming platforms are playing a crucial role in enabling these trends by providing real-time data processing and analytics capabilities. The ability to process and analyze data in real time is essential for developing innovative insurance products, improving customer engagement, and optimizing operational efficiency. As the insurance industry continues to evolve, the adoption of data streaming technologies will be crucial in maintaining competitiveness and driving innovation.

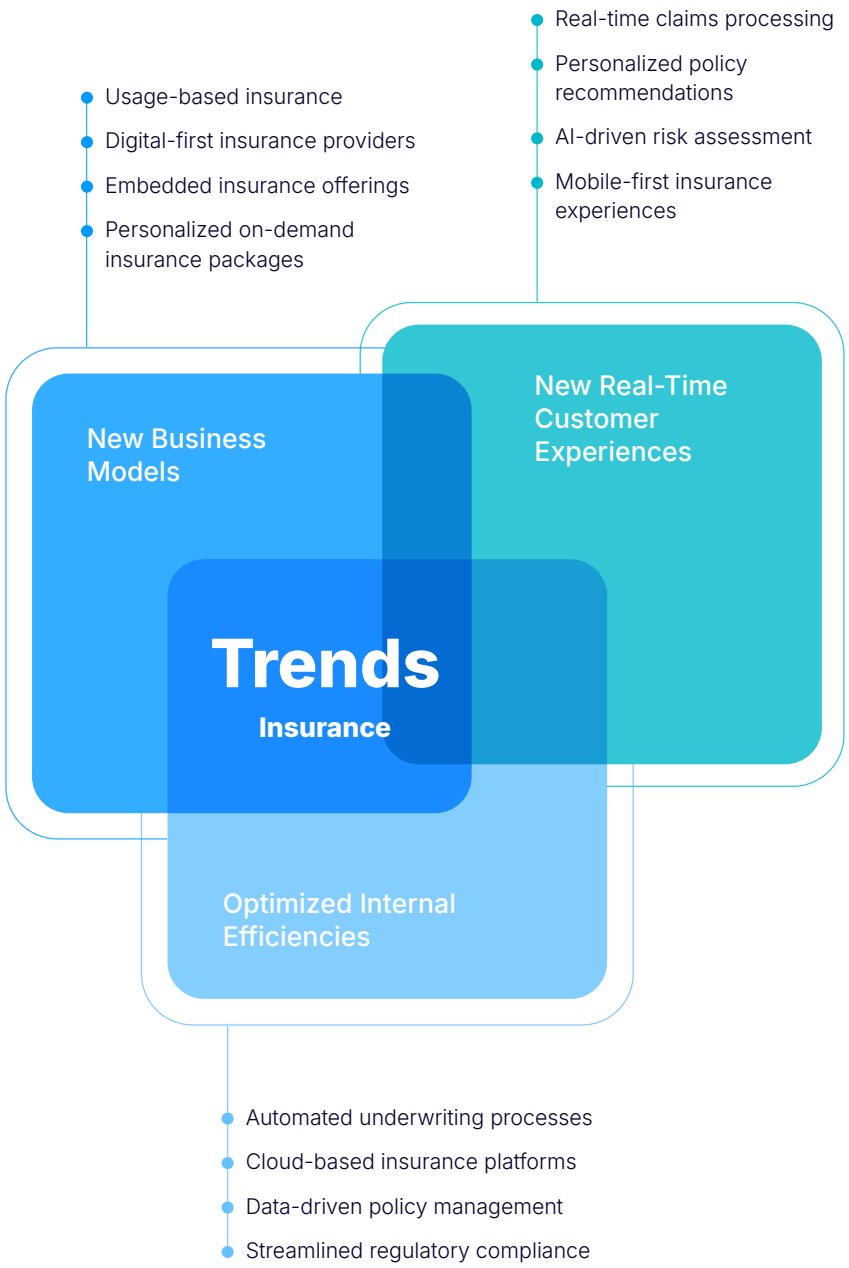

Scenario 1:
Workflow Orchestration

Workflow orchestration in the insurance industry involves automating and coordinating various processes, such as claim processing, policy management, and customer service. Insurers can streamline their workflows by using real-time data from various sources, including customer interactions, policy details, and external data providers. Data streaming ensures data consistency and accuracy, allowing insurers to automate large portions of their processing workflows. More efficient operations lead to more accuracy, ensuring compliance with policy terms and improved customer satisfaction.

Business Value

- **Streamlined workflow automation:** Real-time data processing enables the automation and coordination of various insurance processes, such as claim processing and policy management, improving overall workflow efficiency.
- **Improved data consistency and accuracy:** Continuous data validation and processing ensure data consistency and accuracy, reducing errors and enhancing the reliability of insurance operations.
- **Enhanced claim processing:** Automated review and validation of claims allow for accurate and timely claim processing, ensuring compliance with policy terms and improving customer satisfaction.
- **Increased operational efficiency:** Transitioning from legacy batch processing to real-time workflows reduces processing times and operational costs, enhancing overall efficiency.
- **Comprehensive operational visibility:** Access to real-time data from customer interactions, policy details, and external data providers provides a comprehensive view of claim processing operations, supporting informed decision-making.
- **Improved customer satisfaction:** Efficient and accurate claim processing enhances customer satisfaction by providing timely and reliable service, fostering customer loyalty.
- **Compliance and risk management:** Real-time insights into insurance operations ensure compliance with industry regulations and standards, reducing the risk of penalties and enhancing risk management.

Scenario 2:
Full-Stack Contract Handling

In the insurance industry, full-stack contract handling means managing insurance policies from start to finish. Advanced applications and AI, bring a direct-to-consumer approach to life insurance, speeding up the signing process activating policies within minutes.

Real-time data from various sources like health and driving records offers comprehensive insights of signing and policy activation processes. Data streaming can automate the signing process by quickly contacting third-party data providers and gathering necessary information with the customer's permission. This significantly streamlines the customer experience, cuts operational costs, and improves customer satisfaction.

Business Value

- **Faster underwriting process:** Real-time data processing and automation streamline the underwriting process, reducing the time needed to assess and approve applications.
- **Quick policy activation:** Policies can be activated within minutes, enhancing the customer experience by providing immediate coverage and service.
- **Improved customer satisfaction:** A seamless and efficient process from application to activation leads to higher customer satisfaction and loyalty.
- **Direct-to-consumer approach:** Offering a direct-to-consumer model simplifies the purchasing process, making it more accessible and appealing to customers.
- **Comprehensive data utilization:** Access to real-time data from various sources, such as health and driving records, provides a complete view of the underwriting process, improving decision-making accuracy.
- **Competitive advantage:** Companies that implement full-stack contract handling can differentiate themselves by offering faster, more efficient services, gaining a competitive edge in the market.

Scenario 3:
Micro Insurance

Micro insurance caters to specific, short-term needs, such as insuring brief excursions with rental bikes or similar activities. Advanced applications and IoT technologies enable insurers to implement dynamic pricing models that adjust premiums based on a combination of internal CRM data and external factors. Real-time data analysis made possible by data streaming allows swift adjustments to pricing models to reflect current conditions and individual risk profiles. Through that, insurers are able to optimize their pricing strategies and maintain a competitive edge in the market.

Business Value

- **Tailored insurance solutions:** Real-time data processing allows insurers to offer personalized, risk-based insurance products that cater to specific, short-term needs, enhancing customer satisfaction.
- **Dynamic pricing models:** The ability to adjust premiums based on real-time data, including CRM insights and external factors like weather conditions, ensures competitive and fair pricing.
- **Enhanced customer experience:** By providing insurance solutions that are tailored to individual needs and circumstances, insurers can foster a more engaging and satisfying customer experience.
- **Agility and responsiveness:** Insurers can swiftly respond to changing conditions and customer behaviors, ensuring their offerings remain relevant and attractive.
- **Competitive edge:** The ability to optimize pricing strategies in real time helps insurers maintain a competitive advantage in the market.
- **Increased customer loyalty:** Personalized and responsive insurance offerings can lead to higher customer retention and loyalty.

Real-Life Examples

- **Generali Switzerland**

"At Vitality Group, we also had the opportunity to build from the ground up so there was no limitation on legacy choices. By using Confluent, we found that it also alleviates the day-to-day admin burden. And this basically frees us up with time and allows us to optimize to build our continuous integration (CI)/continuous delivery (CD) build pipeline. And Confluent also allows for seamless scalability so there is no downtime incurred whenever we need to scale."

Ryan James, *Chief Data Officer, Vitality*

Real-Life Example Generali Switzerland

Use Case(s): IT Modernization | Hybrid and Multi-Cloud | Data Governance

Company Overview

Generali Switzerland is a key part of the global Generali Group, one of the world's largest insurance providers. Known for its commitment to innovation, Generali Switzerland offers a broad range of insurance products and services, including life, health, property, and casualty insurance, to both individual and corporate clients. The company has continuously adopted new technologies to improve its operations and deliver exceptional customer experiences.

IT Modernization Initiative

As the insurance industry evolves, Generali Switzerland recognized the need to modernize its core insurance systems to stay competitive and meet rising customer expectations. The company's essential functions—underwriting, policy administration, claims processing, and billing—are critical for assessing risk, managing policies, handling claims, and processing payments. Generali Switzerland embarked on a major IT transformation, using Confluent's data streaming platforms to achieve real-time data processing and analytics, making these functions more responsive and efficient.

The transformation began by deploying a data streaming platform in Generali's data centers, close to its core relational databases. By leveraging real-time data streaming and change data capture, the company could continuously optimize core insurance processes, build new cloud-native applications, and integrate new SaaS solutions. This setup enabled a decoupled, stable architecture, reducing risks while integrating seamlessly with modern platforms that required core insurance data.

Moving to the Cloud

After establishing data streaming in its data centers, Generali Switzerland began migrating the entire platform to the public cloud. This migration was carefully managed to ensure business continuity, with data streaming support via Confluent Replicator enabling a smooth transition from on-premise to a cloud-based data streaming platform. For this first step into a public cloud, Generali moved to managed Kubernetes but retained the self-managed, containerized Confluent Platform. After successful migration and stabilization, Generali continued modernizing by moving to a fully managed Confluent Cloud service, which reduced infrastructure costs and allowed IT teams to focus on business challenges instead of maintaining systems. Again, Confluent Replicator enabled the smooth transition from the ceding Confluent Platform to the new Confluent Cloud.

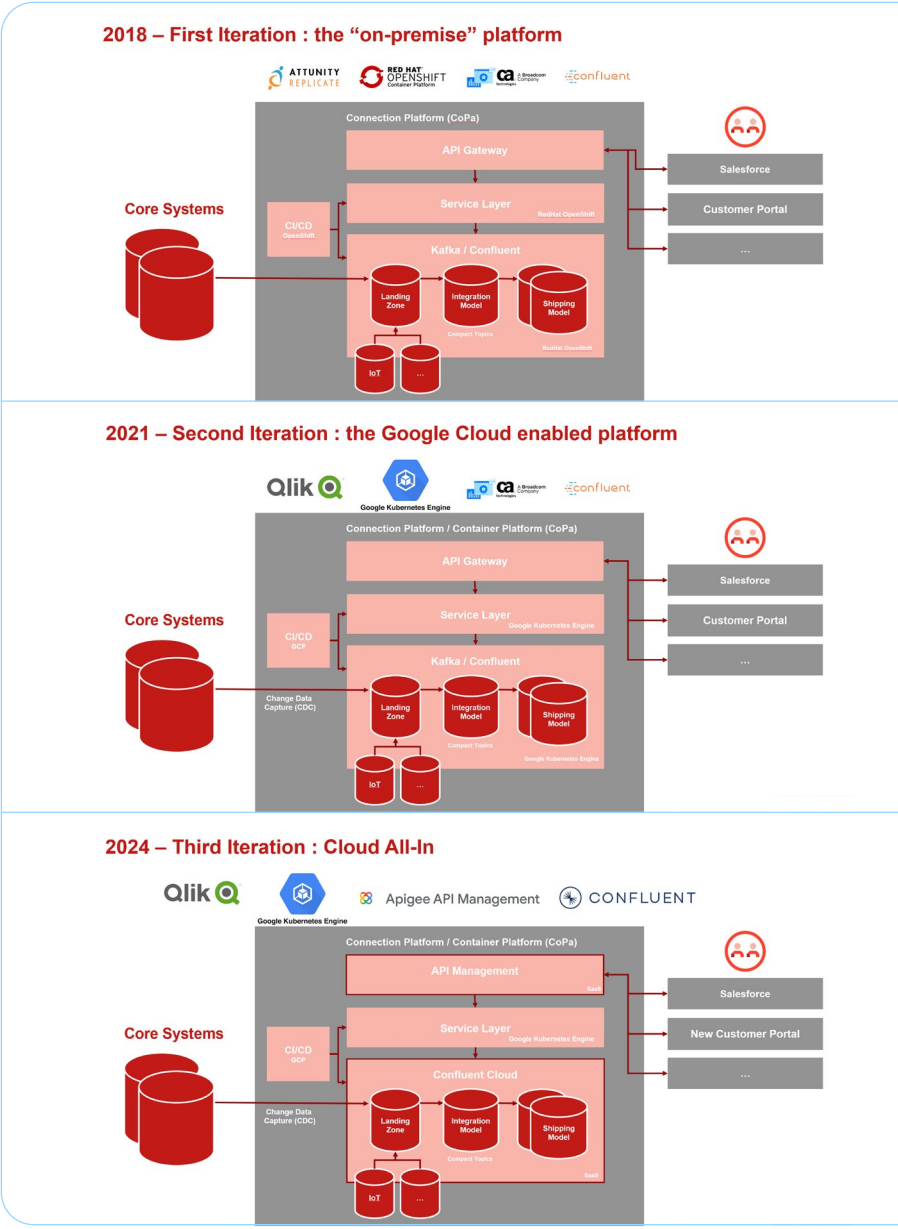

Key Benefits of the Modernization

Generali Switzerland's modernization strategy delivered substantial business value:

- **Speed of innovation for bringing new digital solutions to our customers:** Digital solutions can be integrated and rolled-out faster with minimal risk and minimal to none changes in core systems, thanks to a stable, decoupled architecture. They can either replace legacy core functions, such as a partner master data management system built cloud- and streaming-native or integrated on top of core systems, such as Salesforce Services, Digital Signature or any other Insurtech SaaS services that are plugged into the business process. Decoupled means, that the data streaming layer is decoupling the core and the digital layer from each other.

- **Enhanced operational efficiency:** The enablement of new digital solutions through the streaming layer can increase customer experience, self-servicing, and straight-through processing to a new level which drives operational efficiency. Real-time data processing streamlines core insurance functions, making processes like underwriting, claims handling, and billing faster and more efficient, ultimately boosting operational performance. Integrating legacy databases with modern cloud services has cut operational costs. By reducing the need for maintaining outdated systems, Generali can allocate resources more effectively. New platforms can also be added with minimal risk, thanks to a stable, decoupled architecture.

Leadership Perspective

Generali's Head of IT Alexander Auner states that "our modernization journey has been about building a resilient, decoupled architecture that supports digital transformation without compromising the stability of our core systems such as contract management." For the Lead of the COPA (Connection Platform) Nada Kammoun-Ksibi "moving to Confluent Cloud reduced infrastructure costs and allowed our IT teams to focus on business challenges instead of maintaining systems." "This decoupling has empowered us to integrate high-end APIs and a robust data streaming layer, allowing digital initiatives to thrive while our foundational contract and claims management remains stable and secure," says the Lead Architect Thomas Peter.

> **"The Confluent Cloud is a key driver for our digitalization journey."**

Andreas Schlögl, *CTOO of Generali Switzerland*

Conclusion

The insurance industry is embracing a wave of digital transformation, led by trends such as usage-based insurance, personalized on-demand policies, and embedded offerings. Real-time claims processing, mobile-first experiences, and AI-driven risk assessments are reshaping customer engagement, while insurers optimize operations with automated workflows and data-driven policy management.

Full-stack contract handling illustrates how insurers can deliver fast, seamless service. Continuous data integration from health and driving records ensures quick underwriting and policy activation, often within minutes.
This direct-to-consumer approach simplifies the purchasing process, enhances customer satisfaction, and positions insurers as leaders in convenience and efficiency.

Micro insurance highlights the agility enabled by data streaming. With dynamic pricing models that adjust based on real-time data, insurers can offer short-term, personalized policies tailored to individual risk profiles and external factors. This responsiveness fosters customer trust and loyalty, while providing insurers with a competitive edge.

Real-time data streaming transforms the insurance industry by enabling automation, personalization, and operational optimization. With IoT and AI-driven platforms, insurers can deliver faster, more accurate services while improving compliance and risk management. Use cases such as full-stack contract handling and micro insurance models empower insurers to innovate and adapt to evolving market demands. This shift positions companies to not only enhance customer experiences but also secure long-term growth and leadership in the digital age.

> **"Data is becoming increasingly important and we run everywhere. Data, it's actually money. As an insurer, at the end of the day, we have been always a data essentially industry because our specialty takes really using data in order to price risk. And we are pretty good at it. What now we can do more with data, it's again, what maybe we have not been so good in the past, which is the engagement that we can get with our customer."**

Pietro Carnevale, *Director of Strategy & Innovation, Generali*

3.10 Healthcare and Life Sciences

Debates around digitization are common in the healthcare industry. However, unlike other sectors, these discussions tend to be more cautious, consistent, and understated. One of the most critical trends is the emphasis on data privacy and security, as healthcare organizations handle sensitive patient information. Ensuring the confidentiality, integrity, and availability of this data is paramount, and robust data streaming platforms play a crucial role in achieving this.

Telemedicine and telehealth are also gaining traction, enabling healthcare providers to offer remote consultations and monitoring services. This trend has been accelerated by the COVID-19 pandemic, which highlighted the need for remote healthcare solutions. Remote patient monitoring, facilitated by IoT devices and data streaming technologies, allows healthcare providers to continuously monitor patients' health conditions and intervene when necessary.

The integration of AI in healthcare is another significant trend, enabling advanced diagnostics, personalized treatment plans, and efficient drug discovery processes. Data streaming platforms are essential in enabling these trends by providing real-time data integration and analytics capabilities.

"Physicians need rapid and accurate insights when considering fertility treatments for their patients to ensure timely care. When the models used for each patient are needed promptly, we use Confluent to stream the data from the clinics on-demand. It's as easy as running the request through the model and yielding back a report."

Andrew FigPope, *Lead Data Architect at Celmatix*

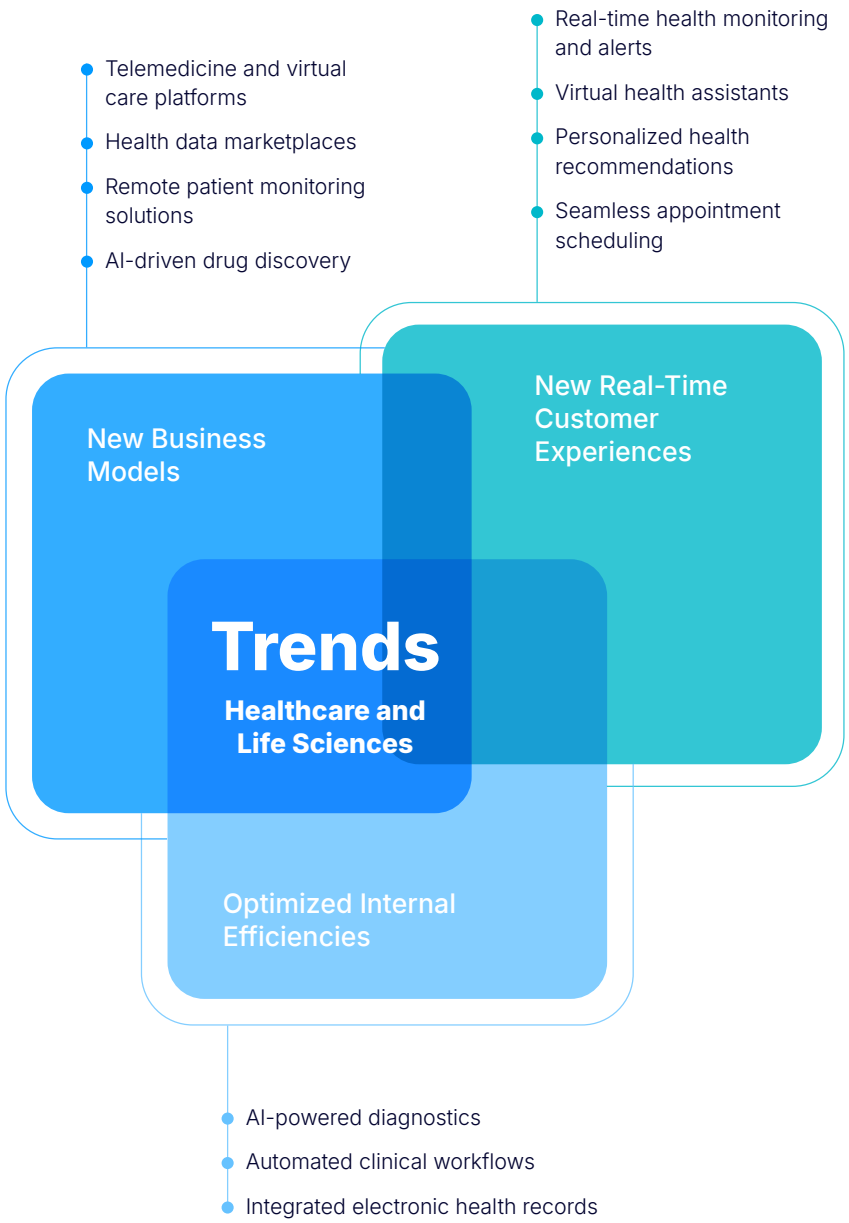

Scenario 1:
Medical Diagnostics

In medical diagnostics, advanced technologies are used to detect and diagnose medical conditions. IoT and AI systems make real-time patient monitoring possible, helping to prevent the risk of deadly conditions. Processing real-time data from various sources like patient monitors, electronic health records, and laboratory results with data streaming technologies enables healthcare providers to monitor patients' health conditions continuously and intervene when necessary, improving patient outcomes and saving lives. A concrete example for the usage of data streaming is monitoring patients that undergo bone marrow transplants for signs of sepsis.

Business Value

- **Improved patient outcomes:** Real-time monitoring and data analysis enable early detection of medical conditions, so timely interventions can improve treatment and save lives.
- **Enhanced patient care:** Continuous monitoring of health conditions enables healthcare providers to offer more personalized care.
- **Early detection and prevention:** Detecting early signs of conditions such as sepsis allows for prompt intervention, reducing the risk of complications.
- **Comprehensive health insights:** Integrating data from patient monitors, electronic health records, and laboratory results offers a holistic view of a patient's health, supporting more accurate diagnoses and treatment plans.
- **Operational efficiency:** Automation of data collection and analysis reduces the need for manual monitoring, freeing up healthcare professionals to focus on direct patient care.
- **Risk management:** Real-time data processing enhances the ability to manage and mitigate risks associated with critical health conditions.
- **Innovation in healthcare:** Leveraging IoT and AI technologies positions healthcare providers at the forefront of medical innovation, offering cutting-edge diagnostic solutions.
- **Cost reduction:** Early detection and intervention can lead to reduced healthcare costs by preventing the escalation of medical conditions and minimizing the need for extensive treatments.

Scenario 2:
Pharmaceutical R&D

Pharmaceutical research and development (R&D) is all about discovery and development of new drugs and treatments. Data streaming and AI/ML technologies enable event-driven workflows and process orchestration, delivering faster data processing to analyze biological images, laboratory results, and clinical trial data more efficiently and with improved data quality. This speeds up the drug discovery process and aids in identifying potential drug candidates, thereby enhancing the efficiency and effectiveness of R&D operations. But while (near) real-time processing is beneficial, maintaining data consistency remains paramount. Data streaming provides consistent access to multiple consumers in the end-to-end data pipeline.

Business Value

- **Accelerated drug discovery:** Continuous data processing and event-driven workflows enable faster analysis of biological images, laboratory results, and clinical trial data, significantly speeding up the drug discovery process.
- **Enhanced R&D efficiency:** By automating data analysis and process orchestration, pharmaceutical companies can improve the efficiency and effectiveness of their R&D operations, reducing time to market for new drugs.
- **Improved decision-making:** Access to consistent information across real-time and non-real-time applications allows researchers to make informed decisions quickly, identifying potential drug candidates and optimizing research strategies.
- **Increased innovation:** Leveraging IT modernization and AI/ML technologies fosters a culture of innovation, enabling the development of cutting-edge treatments and therapies.
- **Cost reduction:** Streamlined processes and efficient data management reduce operational costs associated with drug discovery and development.
- **Comprehensive data insights:** Real-time analysis of diverse data sources provides a holistic view of research progress, supporting more accurate and effective R&D efforts.
- **Regulatory compliance:** Real-time data processing supports compliance with regulatory requirements by ensuring accurate and timely reporting of research findings.

Scenario 3:
Caregiving

Caregiving provides services like senior care and housekeeping. Using data pipelines and applications, caregiving companies can create online marketplaces that connect caregivers with those in need. Real-time data processing capability helps companies offer a wide range of care services, improving customer satisfaction and operational efficiency.

Imagine such an online marketplace using data streaming to match caregivers with clients based on real-time data. This makes the caregiving experience more seamless, leading to more customer satisfaction and business growth.

Business Value

- **Enhanced customer satisfaction:** Real-time data processing enables the delivery of personalized care services, ensuring that individuals receive the right care at the right time, thereby improving customer satisfaction.
- **Efficient caregiver matching:** By utilizing data streaming, caregiving companies can quickly and accurately match caregivers with individuals in need of care, optimizing the allocation of resources and reducing wait times.
- **Comprehensive service offerings:** The ability to process real-time data allows companies to offer a wide range of care services, from senior care to housekeeping, meeting diverse customer needs.
- **Operational efficiency:** Automation of data collection and analysis streamlines caregiving operations, reducing manual intervention and operational costs.
- **Scalability and flexibility:** Data streaming technologies provide the scalability needed to handle large volumes of service requests and adapt to changing market demands, even in the case of extreme situations such as a pandemic.
- **Increased business growth:** By providing a seamless and efficient caregiving experience, companies can attract more customers and expand their market presence.
- **Improved care quality:** Real-time insights into caregiver performance and customer feedback enable continuous improvement of care quality and service delivery.

Conclusion

The healthcare and life sciences sector is advancing rapidly, fueled by trends such as AI-driven drug discovery, telemedicine, and remote patient monitoring. Data streaming technologies revolutionize this landscape, enabling real-time health monitoring, personalized care, and efficient operations while maintaining a steadfast focus on data privacy and security.

In medical diagnostics, continuous data streaming transforms patient monitoring and care. For example, bone marrow transplant patients can be monitored in real time for early signs of sepsis, enabling timely interventions that improve outcomes and reduce complications. This integration of IoT, AI, and real-time data enhances operational efficiency, personalizes care, and positions health-care providers as leaders in medical innovation.

Pharmaceutical R&D leverages continuous data streaming with improved data quality across the data supply chain to accelerate drug discovery. Event-driven workflows analyze biological images, lab results, and clinical trial data more efficiently, reducing cost and time to market for new treatments. These capabilities foster innovation, improve decision-making, and streamline processes. This allows pharmaceutical companies to deliver impactful solutions to patients faster.

Caregiving services further demonstrate the potential of real-time data streaming and data sharing across organizations. By connecting caregivers with clients through online marketplaces, companies offer personalized care and seamless service delivery. Automated processes enhance operational efficiency, while real-time insights into caregiver performance drive continuous quality improvement.

Data streaming enables healthcare organizations to deliver real-time health monitoring, AI-driven diagnostics, and seamless virtual care experiences.

Integrated platforms ensure compliance with stringent regulatory requirements while enhancing patient care and accelerating innovation in life sciences. By adopting real-time data streaming technologies, healthcare and life sciences organizations are equipped to address critical challenges, meet evolving patient expectations, and innovate at scale. These advancements drive better patient outcomes, operational excellence, and groundbreaking research, securing their position as leaders in the digital healthcare revolution.

3.11 Technology

In today's world, technology companies—and increasingly all types of businesses—are transforming into software-defined organizations. This shift, as Confluent's CEO Jay Kreps already stated in 2019, includes that companies move beyond simply using software as a productivity tool. Instead, software now forms the foundation of business operations, with core processes automated, connected, and executed in code. Many of the examples in this chapter could hence be found in previous industry chapters. Technology is reshaping all types of business models, or, as Jay rightly said: Every company is becoming software.

In the technology sector, data streaming has emerged as the critical platform enabling companies to react to data in real time, providing a seamless, instant flow of information across all functions. Especially digital-native companies that are born in the cloud adopt data streaming from the very beginning at the core of their data architectures, running their complete business on a data streaming platform. Several trends are driving this transformation, making real-time data streaming not just beneficial but essential. Cloud computing and Software as a Service (SaaS) offer companies the scale and flexibility to operate globally, while data privacy and cybersecurity remain top priorities amid increasing threats. AI is enabling more dynamic and personalized user experiences, and microservices architectures allow businesses to develop applications with unprecedented agility and scalability.

This chapter explores how data streaming empowers technology companies to not only meet current demands but also build a future-proof foundation for continuous innovation. As companies shift toward event-driven architectures, they unlock the potential to stay competitive, responsive, and resilient in an ever-evolving digital economy. Far from being a passing trend, real-time data streaming is becoming the backbone of the modern enterprise, transforming how businesses operate, engage with customers, and deliver value at scale.

"In terms of scaling with Confluent Cloud, we utilize the possibility to increase or decrease the capacity with a click of a button. It means that when we get ready for Black Friday, or something like that, we can easily scale up the amount of capacity that is required. Also, when we understand that there is lower traffic, we can easily decrease it. So we take advantage of the easy auto-scaling capability of Confluent Cloud."

Natan Silnitsky, *Senior Backend Engineer, Wix*

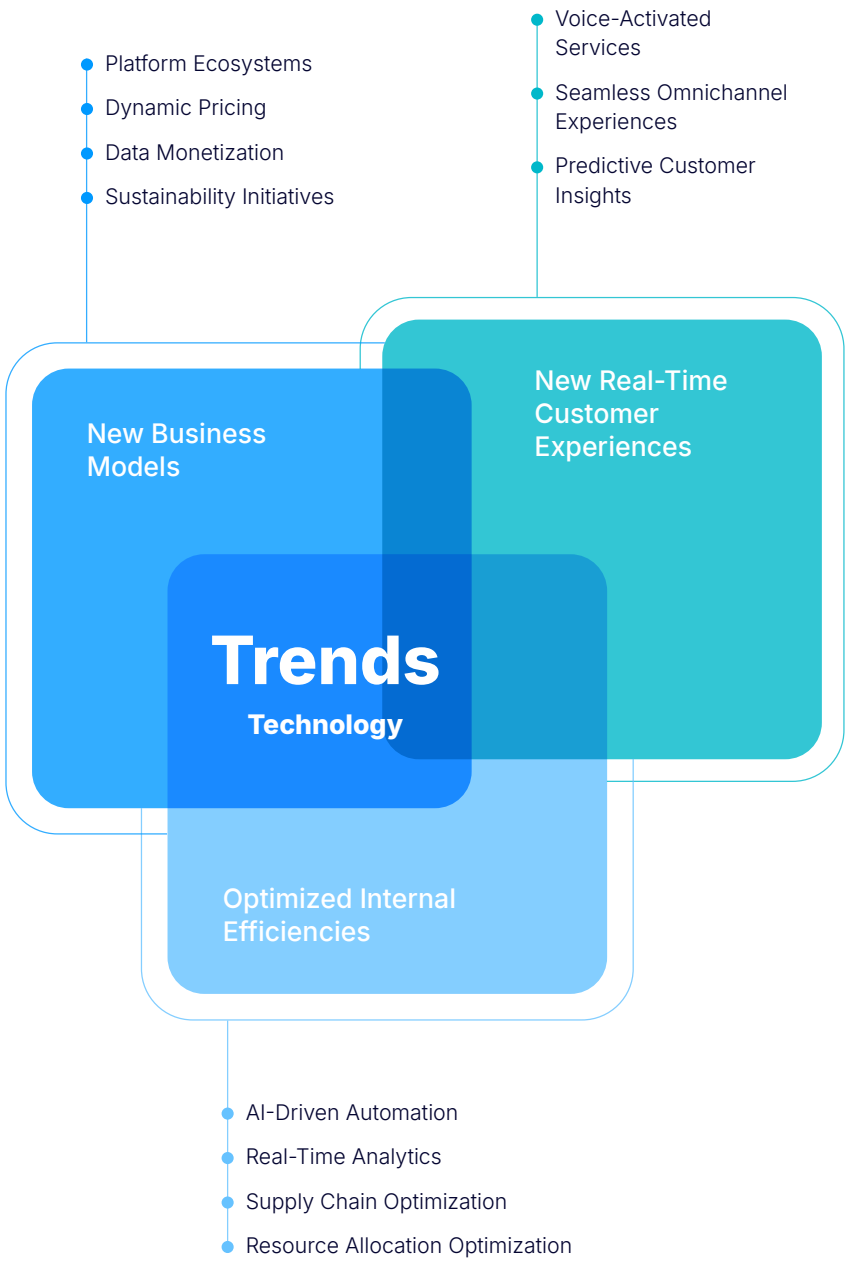

Scenario 1:
Platforms and Mobile Apps

Platforms in the on-demand service or gig economy software category are designed to connect consumers with independent providers for real-time services such as ride-hailing, food delivery, and other logistical services. These applications can be described using several specific categories:

- **On-Demand service platforms:** This is the most common category, as these apps provide services on demand by connecting users to service providers based on real-time needs. Examples include platforms for ride-hailing, grocery delivery, and food delivery.
- **Marketplace platforms:** These platforms function as digital marketplaces where consumers and providers meet to transact. They include features such as payment processing, customer service, and logistics, facilitating smooth transactions between parties.
- **Logistics and delivery software:** For apps that manage inventory, delivery schedules, and driver routes, they also fall under logistics software.
- **Mobility-as-a-service (MaaS):** Transportation services that offer an alternative to vehicle ownership by providing transportation as a flexible service fit within this category.

Data streaming is revolutionizing the way platforms and mobile apps operate, by delivering dynamic content, personalizing user interactions, and ensuring seamless performance across platforms. Real-time data streaming allows for the continuous flow of information, enabling platforms and apps to respond instantly to user actions and external events.

For instance, a social media platform can harness the power of data streaming to deliver real-time updates on user feeds. This capability ensures that users are constantly engaged with the freshest content, whether it's a friend's status update, a trending news article, or a live event broadcast. Equally for an (on-demand) service platform like a ride-sharing app, data streaming allows users to see the exact location of their driver in real time, receive instant updates on estimated arrival times, and get notified of any route changes due to traffic conditions. This real-time interaction not only enhances the user experience by providing transparency and reliability but also builds trust and satisfaction as users feel informed and in control of their service experience.

Similarly, an e-commerce app can use data streaming for personalized product recommendations by analyzing a user's browsing history, purchase behavior, and current market trends. This allows the app to adjust recommendations in real time, for example, to showcase hiking gear for a user interested in outdoor equipment. This level of personalization improves the shopping experience by making it more relevant for the user and boosts conversion rates. Real-time data streaming enables the app to offer timely discounts or suggest complementary products at the moment of consideration.

Business Value

- **Enhanced user engagement:** Real-time updates and dynamic content keep users constantly engaged, increasing the time spent on the platform and fostering a more interactive experience.
- **Personalized user experience:** By analyzing user behavior in real time, platforms can offer personalized interactions and recommendations, enhancing user satisfaction and loyalty.
- **Improved transparency and trust:** Real-time data streaming provides users with up-to-date information, such as driver locations or order statuses, building trust and reliability in the service.
- **Operational efficiency:** Automation of data processing and logistics management reduces manual intervention, streamlining operations and lowering costs.
- **Increased conversion rates:** Personalized recommendations and timely offers based on real-time data analysis can significantly boost conversion rates and sales.
- **Scalability and flexibility:** Data streaming technologies provide the scalability needed to handle large volumes of transactions and user interactions, supporting business growth and expansion.
- **Competitive advantage:** Companies that leverage real-time data streaming can differentiate themselves by offering superior user experiences and innovative features, gaining a competitive edge in the market.
- **Proactive issue resolution:** Continuous monitoring and real-time insights allow for the quick identification and resolution of potential issues, minimizing downtime and maintaining service quality.
- **Informed decision-making:** Access to real-time data enables companies to make informed decisions quickly, optimizing operations and strategic planning.

Scenario 2:
Infrastructure Monitoring

Data streaming is transforming infrastructure monitoring, particularly for companies that require real-time insights for optimal performance and security. By leveraging data streaming, these companies can continuously collect, process, and analyze data from various sources, enabling proactive monitoring and rapid response to issues. This capability ensures that systems remain robust, efficient, and secure, and enhance customer satisfaction.

For example, a cloud service provider can use data streaming to monitor server performance, addressing bottlenecks before they impact users. Similarly, a cybersecurity firm can utilize data streaming to detect anomalies in real time, allowing for immediate intervention to prevent breaches and safeguard sensitive information.

Business Value

- **Proactive issue resolution:** Real-time monitoring allows companies to identify and address potential issues before they escalate, minimizing downtime and maintaining service quality.
- **Enhanced security:** Continuous data analysis enables the rapid detection of anomalies and potential threats, allowing for immediate intervention to protect sensitive information.
- **Operational efficiency:** Automation of monitoring processes reduces the need for manual oversight, lowering operational costs and increasing efficiency.
- **Improved system performance:** By continuously optimizing system performance through real-time insights, companies can ensure high availability and reliability of their services.
- **Scalability:** Data streaming technologies provide the scalability needed to handle large volumes of monitoring data, supporting business growth and expansion.
- **Innovation and agility:** Data streaming fosters a culture of innovation, enabling companies to quickly adapt to changing market demands and technological advancements.

Conclusion

The technology industry is at the forefront of a global shift toward software-defined operations, where real-time data streaming is a critical enabler of innovation, agility, and scalability. Emerging business models such as platform ecosystems, data monetization, and dynamic pricing are reshaping the landscape, while seamless omnichannel experiences and AI-driven automation define better customer expectations. Continuous data processing forms the backbone of these transformations, driving efficiency, innovation, and engagement.

Platforms and mobile apps illustrate the power of real-time data. On-demand services like ride-hailing apps leverage data streaming to provide instant updates on driver locations, route changes, and estimated arrival times to create transparency and trust. Digital platforms use real-time nsights to offer personalized product recommendations and dynamic pricing, ensuring a tailored and engaging shopping experience. These capabilities not only enhance user satisfaction but also significantly improve conversion rates and loyalty.

Infrastructure monitoring showcases how real-time data streaming optimizes internal efficiencies. Service providers, for instance, continuously monitor server performance to address bottlenecks before they impact users. Cybersecurity firms use continuous data streaming to detect anomalies in real time, enabling immediate intervention and robust data protection. These applications ensure high availability, operational efficiency, and strengthened security to provide a competitive edge in a fast-paced environment.

Continuous data streaming also supports sustainability initiatives by optimizing resource allocation and energy use. For example, data-driven insights into supply chains allow for better planning and reduced waste, aligning with global sustainability goals. Companies leveraging cloud-native data streaming technologies can scale operations flexibly, adapt to market changes swiftly, and deliver continuous innovation.

As companies adopt event-driven architectures, they unlock unprecedented opportunities to innovate, personalize, and scale. Data streaming empowers technology companies to enhance user engagement with predictive insights, ensure operational resilience through real-time monitoring, and innovate rapidly in response to evolving market demands. By embedding real-time data capabilities into their operations, businesses can transform into agile, future-ready enterprises that thrive in an increasingly digital economy.

Creation and Positioning of a Data Streaming Organization and Community

04

Introducing a data streaming platform is more than just a technical shift—it requires a fundamental change in the business. As such, it must be managed like any other transformational initiative, aligning with key change management principles. This includes strategic internal communication and leadership buy-in, requiring awareness of the benefits of data streaming both vertically and horizontally within the organization. Community building is crucial in this process, serving as a hub for exchanging use cases and successes. A vibrant community not only educates and inspires but also ensures the platform's continued relevance and usage.
Key change management principles supporting this approach include:

- **Communication:** Ensuring consistent, clear messaging across the organization to build awareness and understanding of the platform.
- **Leadership Alignment:** Securing executive sponsorship to champion the initiative, ensure resourcing and guide teams through the transition.
- **Engagement and Involvement:** Involving all relevant teams early in the process, fostering cross-functional collaboration, and encouraging shared ownership.
- **Managing Resistance:** Addressing potential resistance by timing the community's launch carefully and providing the right support structures.
- **Education and Training:** Equipping stakeholders with the knowledge and skills they need to adopt and implement the platform effectively.
- **Sustaining Change:** Building momentum through ongoing communication and showcasing successful use cases to sustain engagement over time.

Framing the adoption of a data streaming platform through these principles helps to make the platform an integral part of operations, driving innovation and competitive advantage. An active data streaming community is the core of the Data Streaming Organization (DSO), a framework designed to guide organizations to meet increasingly complex customer needs with governed, shared, accessible, and real-time data. This is achieved by providing a unified technology platform, established ways of working, and a guiding data streaming strategy and vision.

If you want to read more about the holistic framework of the DSO, we highly recommend *the whitepaper "Maximize Data Value and Innovation Across the Organization"* written by the Confluent Strategic Advisory Team.
In previous chapters, we shared several use cases that were first introduced within internal data streaming communities. Sharing the success stories, companies were able to encourage teams to achieve their business goals with data streaming. Successful use cases should be highlighted through various channels, which we will discuss in this chapter.

This book is one such channel for showcasing use cases. If you believe your use case would be a good fit to be featured in this book, feel free to reach out to us through ✉ use-case-book@confluent.io.

4.1 The Journey to a Data Streaming Organization

Data streaming revolutionizes how companies process and use their data. A data streaming platform can be a powerful component of any data architecture—as long as it is used and communicated correctly. For this to happen, it needs the buy-in of all stakeholders involved, from business to IT. While clear internal positioning and consistent, continuous, and emotionalized communication are all essential factors, they are all second to one crucial requirement: timing. Starting a community can easily do more harm than good if it does not happen at the right time with the right message. Companies need to have reached a certain level of maturity in data streaming before any communication strategy around can be put into place.

Back in 2018, Lyndon Hedderly set out a maturity model in his white paper *"Five Stages to Streaming Platform Adoption"* that has lost none of its relevance: his guide describes the enterprise adoption process of Apache Kafka and a data streaming platform in five stages, from awareness-building all the way to full integration of the platform as the central nervous system of the company. The direction of the initial drive within the company varies greatly, but the "bottom-up" approach (initiation comes from the project team) appears to be more common than the "top-down" approach (initiative is started by management).

The "Enterprise Adoption Barrier"—as displayed in the next graphic—is a frequent phenomenon that occurs within organizations as they move up the maturity curve following the bottom-up approach. Moving from stage 3 to 4 often requires a disruptive organizational and processual transformation—without it, companies hit the aforementioned barrier, effectively hindering further adoption.

Another indispensable prerequisite is a centralized team for enterprise services, even with a fully managed service like Confluent Cloud. Someone in the organization still needs to define and own governance, cost accounting/attribution, and a common operating model for infrastructure, while offering user-friendly self-service tools and APIs. These challenges often lead through a somewhat painful process of developing centralized coordination, strategy, and oversight after data streaming has been widely adopted.

We want to briefly highlight the individual stages of data streaming adoption, based on the most common bottom-up approach. You will find that these phases often reflect the development steps and maturity stages of individual business units within a company as well as the entire organization.

Unified Data Streaming Strategy

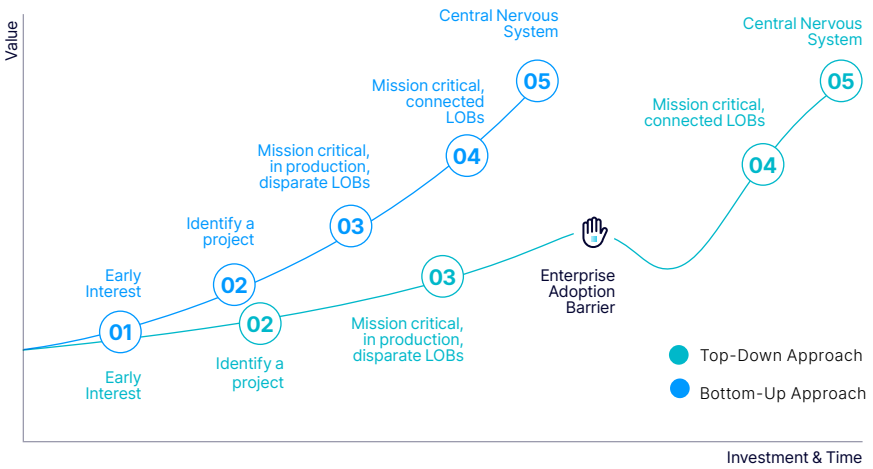

In the data streaming maturity model, the difference between a top-down and a bottom-up approach lays in the enterprise adoption barrier

Stage 1: Early Interest—pilot projects and initial successes

The journey often begins with small pilot projects initiated by technology teams, running non-critical use cases. This phase makes it possible to test initial applications and experience the agility and flexibility of a data streaming platform. Teams usually benefit from increased development speed and the new possibilities of real-time data processing.

Stage 2: Identify a project—production-ready applications

The company is starting to implement the first applications in production. These applications offer both technical and business benefits, such as improved customer experiences and the development of new business models. This phase marks the transition from experimental applications to mission-critical systems.

Stage 3: Mission critical, in production disparate LOBs—integrated streaming and strategic alignment

In this phase, the data streaming platform is closely linked to the company's strategic goals. Different business units work together to integrate data streams and break down silos. The platform supports the security, reliability, and scalability of data processing. The establishment of a central streaming service and/or a Center of Excellence (CoE) is usually initiated in this phase with the clear goal to move to stage 4. Decisions made in this stage include core components of a

Data Streaming Organization which also includes the establishment of an internal community around data streaming and creating an integrated communication and positioning concept surrounding the central service.

Stage 4: Mission critical, in production connected LOBs—Global streaming and cloud integration

In the global streaming stage, the platform is scaled across the enterprise. The company starts using the advantages of cloud services to process data efficiently and comply with more legal requirements. The platform enables the company to make data available in different geographical regions and consolidate global business processes at the same time. The internal community is established and growing, building upon a multi-channel communication strategy that ensures the continued buzz about successes and future strategies.

Stage 5: The central nervous system of the company

In the final stage, companies become true Data Streaming Organizations, where the data streaming platform becomes the central nervous system. All data is available in real time, which significantly improves efficiency and the ability to innovate. Companies at this stage demonstrate the power of this architecture by processing large amounts of data in real time and continuously developing new applications.

"Companies that invest more in data streaming also tend to have more mature implementations. 68% of organizations in Level 4 cite data streaming as a top strategic priority, compared to 48% in Level 3—signaling a strengthening commitment to data streaming as organizations mature in its use."

Data Streaming Report 2023

Transitioning an entire organization to produce and consume streaming data usually requires significant training and education to promote this new technology paradigm. There is always at least some technical infrastructure gap that necessitates homebrew solutioning. This can become an expensive endeavor and requires executive sponsorship to some extent. Once set up, other parts of the business must align—something that comes with the significant burden of cross-functional execution.

Check out the whitepaper on how to create a Data Streaming Organization

4.2 Key Drivers for a Community Within a Data Streaming Organization

Creating a unified, reliable source of information for a particular technology often demands internal support and persuasion. Typical questions, such as "Is the platform/service ready?", "Is this topic important for our goals?", or "Do we have meaningful results to share yet?" are natural and should be confidently answered with a "Yes!" by the project team before getting started.

The motivations behind forming a communication strategy and community are as follows:

- **Internal (and external) positioning**
- **Know-how transfer**
- **Consolidation of projects and use of best practices**
- **Hiring (internal & external)**

Internal (and external) positioning

Apache Kafka, Apache Flink, and Confluent are innovative technologies that open up new opportunities for companies. Project teams, owning the central data streaming service, get this and believe in the strategic relevance the paradigm of data streaming brings for the organization. To have this paradigm understood across the organization, the concept, the vision, and the benefits need to be positioned and seen.

So it's not just about keeping your position in your market by executing a successful data strategy. It is also about positioning the internal data streaming service within the company in order to reach the next stage on the streaming maturity curve and create more added value through the platform. That means positioning the CoE and project team itself within the organization. The message should be clear, and it should be confident: This department makes a significant, if not *THE* difference in the company-wide digitization strategy! Don't hide your light under a bushel!

Regardless if your data streaming strategy is initiated through a "bottom-up" or a "top-down" approach: In both cases, the opposite party should be involved as quickly as possible, because no strategy works without implementation and vice versa (also see "Enterprise Adoption Barrier"). To achieve this, the strategic relevance, the actual functionalities, and the vision of the platform must be communicated all across the enterprise and, most importantly, to all different target groups.

Equally important as the internal communication is the external visibility. Because, spoiler alert: For some colleagues, the penny will only drop when they read about the company's internal data streaming project—to put it bluntly—in the daily newspaper over breakfast. We will see concrete examples for internal and external positioning later in this chapter.

Good news for the "bottom-up" project teams: **For 51% of IT leaders** surveyed in the *2024 Data Streaming Report*, data streaming is one of the top strategic priorities for IT investment in 2024, compared to 44% in 2023. In addition, **68% expect the use of this technology to continue to increase over the next two years.** But even if the adoption of data streaming increases, it is important to continuously engage all stakeholders, both through internal and external visibility and activation.

Know-how transfer

The success of any technology implementation depends on how well it is understood internally and how effectively that knowledge is shared. Many service owners and CoEs are therefore in close partnership with the internal training or academy departments. These departments usually have the necessary structures in place to make training opportunities visible and accessible within the company.

> **"The decisive thing about knowledge is that you take it to heart and apply it."**
>
> *Confucius*

Many companies rely on the free on-demand video courses by Confluent for the basics on Kafka, Flink, and beyond, which offer a good introduction for developers and architects in videos that build upon each other.

In addition to building up technical knowledge, be careful not to lose sight of the conceptual and, above all, strategic aspects. For that, consider running a variety of educational sessions where you explain the core concepts of data streaming and an event-driven architecture to non-IT departments and stakeholders.

Consolidation of projects and use of best practices

Open source is a wonderful thing! Anyone can use it and test it out with new use cases. As wonderful as this may seem, it is sometimes a thorn in the side of many IT departments. Some may even hush conspiratorially about the infamous and unpleasant thing called "shadow IT." This makes it even more important for service owners to create a central point of contact for all matters in this area as quickly as possible. In other words: You want Kafka? We can do Kafka. And we can do Kafka in a structured and compliant way to your benefit! A central data streaming platform only makes sense when it is actually central and the whole company can publish and subscribe to its streams of data in a governed way.

Besides that, we noticed several occurrences with clients where different teams were working on similar use cases behind office door 1 as well as behind office door 3. The time saved if those teams met through an internal community is obvious. Over time, there is no need to reinvent the wheel, as the internal community also serves to distribute architectural blueprints and give guidance on regulations and compliance.

It is recommended to check with your internal communication team what the most common collaboration platforms and channels are, so you can choose an environment that your colleagues feel most at home with.

Hiring (internal & external)

We will explore the options for presenting the project externally in order to attract new talent (so called Employer Branding) later in this chapter. However, internal recruitment shouldn´t be ignored either. Successful projects still require a dedicated team for ongoing development, even in a Software-as-a-Service world. Often, reallocating internal resources is the faster route to accelerate project growth. When the strategic relevance of the project is well-known within the company and potential career opportunities are clear, internal interest tends to follow quickly. In fact, some of the largest data streaming services among our clients started from internal demand and interest, allowing teams to scale rapidly.

Executive sponsorship is an absolute accelerator for internal team expansion, but it must be earned through a clear articulation of the platform's mission and vision, alongside demonstrated initial successes and a professional presentation backed by a consistent branding strategy.

4.3 Fundamental Tactics for Developing a Data Streaming Community

After kicking off this chapter with change management principles to successfully establish a Data Streaming Organization, we will now combine them with the core aspects of marketing like personas, channels, and promotion.

Addressing your Target Group

Anyone who has ever enthusiastically talked about his or her Kafka project at the (virtual) coffee machine will be familiar with the response "Oh, you're into literature?" Even though Apache Kafka has become the global de facto standard for data streaming, non-IT people usually associate Kafka with the author and data streaming with the latest Netflix series (which is not entirely wrong, Netflix runs its recommendation engine on Kafka). When talking about the advantages of data streaming technology, you cannot rely on one common understanding for everybody.

That is why it is so crucial to employ different tonality, messaging, and forms of presentation for each persona, like senior management, C-levels, architects, application owners, and developers. Each one of those roles speak a different language and focus on different aspects of data streaming.

While developing these enablement and training plans, you should always keep in mind that everybody learns differently. This fact must be taken into account if you want get everyone in the company to buy in. The larger the company, the longer it will take to gain a basic understanding of what a data streaming platform and terms like "event-driven" actually mean and what opportunities they offer for IT and business units. This also means that we need patience in combination with a continuous communication strategy on all channels and across all learning styles:

- **Visual learning**
- **Auditory learning**
- **Reading and writing**
- **Kinesthetic learning**

It is recommended to provide material for as many learning styles as possible, both stand-alone assets (videos, podcasts, blog articles) and perpetually accessible information in a single source of truth. Dedicated "Kafka Days" or "Data Streaming Days" are becoming increasingly popular at Conluent and offer the opportunity to convey the technology and added value via a wide variety of channels within a single day. We will cover some examples later in this chapter.

There might be dedicated personas or departments to target first, like solution architect groups or analytical teams. It is important to understand that even the "low hanging fruits" in terms of use case recruitment will be at different stages of maturity themselves. Although the company may already be at level 3 or 4 of the maturity curve, you will always come across departments/ business units/application teams that are still at level 1 or 2. This requires a different approach compared to a team that has been working with the technology for a long time. This is another reason why a continuous and consistent communication strategy with one easily accessible single source of truth is indispensable.

After analyzing your target group and the corresponding personas, you can move on to the next step: **How and where do we reach them?**

Finding Your Communication Channels

One goal of a community within a data streaming organization is to proactively prevent "reinventing the wheel." The same principle applies to creating the community itself: many companies already have well-functioning communities (often role-based, such as architecture tribes or IT management roundtables, or topic-based, like those focused on GenAI or microservices). It's worth exploring best practices internally: How were successful communities started? Are existing groups willing to support the new one and offer a virtual handshake? What channels were used to establish the current communities, and who manages these? Where do these communities "live"? Familiarize yourself with the internal culture of collaboration. The structure and setup of these channels can be as diverse as the companies itself, but a good starting point is to reach out to the internal communications department. You may even have a dedicated team for community management or employee experience. They are often eager to get involved early on, as this allows you to align the community's setup with their standards and leverage their existing resources and frameworks. Plus, having this department as an ally is always a strategic advantage.

Examples of common internal channels to promote your new community and its activities:

- Existing community channels where the target groups is already engaged
- Internal newsletters
- Dedicated mailings to target group
- Existing educational formats (lunch & learns, town hall meetings, ...)
- Promotional screens in the office (yes, the ones that also show you the daily menu of the cafe)
- Posters

Building your Message

Once we've identified the "who" and "where," the next, more labor-intensive step is tackling the "how." As mentioned earlier, different individuals have varied learning styles and respond to different communication approaches depending on their roles. The challenge now is to distill these differences into a common denominator, which is often (sorry, tech enthusiasts) less about technical details and more about delivering a strategic and emotional message. This message needs to be consistently and continuously communicated through the identified channels.

By aligning the core messages with your company's (technology) mission and vision ensures that the project's value is recognized within the company and positively associated with the corporate strategy. Using a familiar color palette and visual language can also enhance recognition.

Based on our experience, we recommend a multi-phase launch of the internal community. Start by addressing the broad target audience, then refine the messaging to fit specific personas, conveying the added value in a way that resonates with each group.

Make sure to regularly position yourselves with new and growing projects, either within the company's existing internal formats or by establishing new ones. In other words, "Do good and talk about it!"

Figuring out the "who," where," and "how" is not an easy task. Confluent offers a dedicated workshop on the topic of "Positioning & Branding."

Reach out to the Strategic Advisory team at Confluent to learn more about Community Building and the DSO.

Examples of internal community branding at various companies

4.4 Strategies for Long-Term Engagement in a Data Streaming Community

From this point on, it is important to continuously meet the needs of all personas across the identified channels. If successful, this step will be repeated several times, especially in the first few months (or years, depending on the size of the company), as new teams continuously bring their use cases to the central data streaming platform. Different levels of knowledge and expectations will be met along the journey and it is recommended to meet them proactively instead of having to be reactive.

A popular approach to kick off and maintain a positioning strategy are internal events. We would like to showcase two examples that can also easily be combined.

Activities to Build a Community

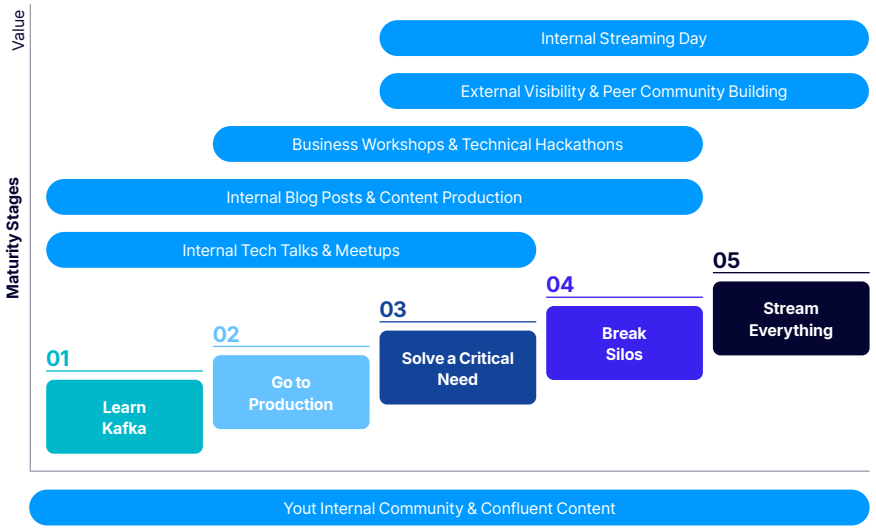

Examples of community engagement programs, along the internal maturity curve

Kick Off Session & Series of Educational Talks

With one initial session, we want to make sure we get everyone at the same level of understanding what data streaming actually is. Hence we recommend a 60-minute session on the topic "What is data streaming? Typical use cases in our industry."
It is important that the welcoming or opening notes to this presentation are not done by the technology provider, but rather by the internal project team who introduces themselves. At best, the mission and vision of the internal service can be communicated within the first 5-10 minutes. Also make sure to introduce the channels where interested teams can reach out to you (a simple email alias can do the trick) and where they can find more information.

The important part is to not only introduce static informative sources and collateral, but also future educational events, office hours, and/or interactive Q&A formats. The teams need to understand that what you are introducing really is a strategic initiative that is highly relevant for the company's digital future.

Also, make sure to include an opportunity for interested teams to share their first successes and use cases: Be as open and collaborative as possible and make sure all output is aligned with the internal branding that you as a project team / service owner team defined.

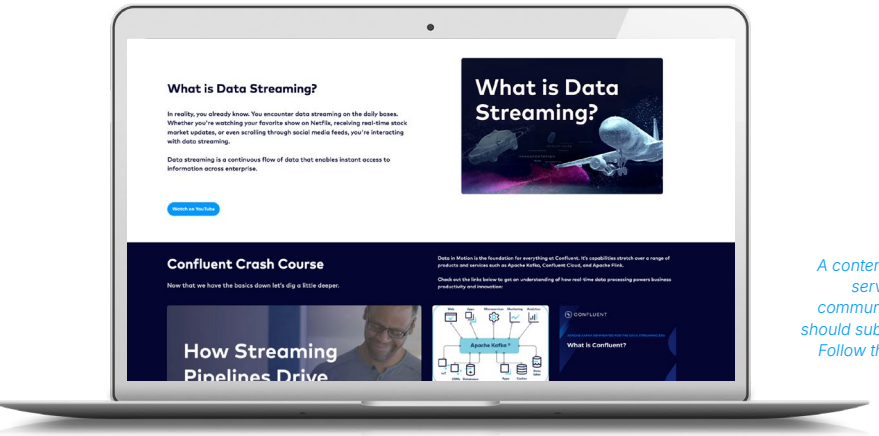

A content delivery service every community owner should subscribe to: Follow the Stream

Full Day Event: Data Streaming Day / Kafka Day

Depending on your company's culture and its position on the maturity curve, hosting a dedicated "Data Streaming Day" or "Kafka Day" might be a smarter, more impactful way to raise internal awareness. This approach is especially effective if several use cases are already live and ready to be showcased, or if there's high interest in certain hot topics (e.g., GenAI, customer 360, digital twin, etc.).

The beauty of such an event is that you can leverage all types of event management strategies. Here are some ideas to make the most out of it:

- **Dedicated branding** for the event, ideally co-branded with the community, and personalized either regionally or thematically (e.g., BMW's "Kafka Wiesn" under the slogan "O'gstreamt is," or s.Oliver's "Kafka Fashion Convention" with the tagline "Fashion is a Stream of Events").
- **Promotion materials** like posters, flyers, digital slides in the office, and e-mail newsletters to create buzz and anticipation.
- **Creative giveaways** for participants, such as community-branded T-shirts, stickers, or even customized cupcakes—get creative and make it fun and tasty!
- **Recordings and slide decks** from the event presentations to share in your community channels afterwards and repurpose it for further content creation.

BMW Kafka Wiesn (Internal Oktoberfest)

The BMW Kafka Wiesn celebrated its 5-year anniversary in 2024, bringing together the global BMW Group Kafka community in a hybrid event in the spirit of Oktoberfest and innovation. Every year, BMW Group showcases their most recent use cases, strategic initiatives, and the vision of their internal data streaming platform. One thing that cannot be missed: The Kafka Pretzels!

Michelin Kafka Day

At the Michelin Kafka Day in May 2024, Michelin's global data streaming community was committed to the topic and the vision. The agenda included a variety of formats, which were offered to over 500 participants across three time zones.
In addition to Michelin-branded macarons, the exclusive event shirts also sold like hot cakes

Regardless of the format in which the data streaming platform and the associated community are introduced: Everything has to happen on-brand and via the channels of the respective community. Continuity and permanence are essential, especially at the beginning, so that an internal brand can be built up. Another aspect to watch out for is to provide internal content and knowledge by the technology provider in equal measures—or the recipients might gain the impression of a "one-sided advertising event," which can further lead to a negative bias towards the technology vendor and, in turn, toward the product. It is better to act as a joint team right from the start, with the focus on internal use cases. If new topics are to be presented, it is better to directly relate them to internal projects, strategies, or visions.

External Engagement to Strengthen Brand and Community Recognition

When it comes to positioning and perception, the world beyond your company should not be overlooked. Once the first use cases have been successfully implemented on your data streaming platform and shared within the internal community, it's time to extend that communication externally. But what does that look like in practical terms?

1. **Become part of the Confluent community and share your successes and experiences with other customers.** This fosters a valuable exchange with industry peers or others working on similar use cases. Why reinvent the wheel when others might have useful insights? Confluent offers various formats for this, including one-on-one phone exchanges, guest speaker spots at customer internal Kafka Days, and industry-focused roundtables. Want to connect with multiple customers in your region at once? The Data in Motion Tour and the global data streaming conference Current provide excellent platforms to engage with customers across different industries and maturity levels.

2. **Apache Kafka** and **Apache Flink** being open-source communities also offers great opportunities to showcase technical expertise, share experiences, and receive feedback. Over 50,000 members meet regularly in over 260 Meetups across 66 countries—and they're always looking for new hosts and speakers.

3. **Your success stories deserve to be shared with the world**—whether digitally or in print. Formats like YouTube videos, written case studies, webinars, and press interviews provide ample opportunities to boost brand awareness and employer branding. After all, who wouldn't want to work for a company that's driving innovation with cutting-edge technology?

The beauty of these external formats is not only in the positive presentation of your project to the outside world but also in the new content they generate to further raise awareness internally for your community, service, and team.

Examples of external presence and positioning include:

- ***Video of Michelin*** showcasing their gold standard reputation for innovative mobility solutions
- ***Speaking at Current***, the biggest event in data streaming
- ***Participating in a Webinar*** alongside guest speakers and Confluent experts
- **Being featured on the *Confluent Customer Hub***, surrounded by peers and industry leaders

Data Streaming Is a Journey: A Roadmap to Success

05

Data streaming is not just a technological advancement; it is a journey. This journey complements other data management technologies like data lakes, data warehouses, and lakehouses. Embracing data streaming requires a shift in mindset, moving from traditional data processing methods with batch workloads and point-to-point APIs to a more continuous, real-time approach with a decoupled, flexible event-driven enterprise architecture. Every journey starts with a first step so we hope that this book gives you the courage to either start your journey or to accelerate your existing one to expand the usage across your organization.

Technical Data Streaming Use Cases and Industry Success Stories

This book has explored various technical use cases across many industries. While the core principles of data streaming remain the same, the value it adds can vary greatly depending on the specific scenarios and industry needs. Whether it's improving personalized customer experiences in retail, optimizing production lines in manufacturing, or enabling fraud prevention in financial services, data streaming offers a wide range of benefits tailored to each industry's and business units unique requirements.

Industries are always changing, driven by new trends and consumer demands. To thrive in this dynamic environment, organizations need agile and adaptable architectures that enable them to seize opportunities and tackle challenges head-on—and finally stay ahead of the curve.

As transformative technologies like AI emerge, they bring not only opportunities but also a mandate to rethink strategies. By embedding data streaming at the heart of your modern data architecture, companies unlock its immense potential to drive innovation, foster growth, and redefine the competitive edge. This book serves as a guide to that journey, offering insights, trends, and technical applications that will empower organizations across all industries to harness data streaming's full capabilities. Hence it is also the plan to constantly extend this book with new trends, technical use cases, and, most importantly, new industry examples and success stories. This resource aims to remain dynamic and relevant, reflecting the latest advancements and insights in the field.

Build a Unified Data Streaming Ecosystem Through Community

Following the five-stage maturity curve, organizations evolve their data streaming initiatives from small pilot projects to becoming the central nervous system of the enterprise. The path to success with data streaming holds the plan for an unified platform, a strategic roadmap, and a focus on fostering an internal community. Data streaming platforms thrive when supported by a strong Data Streaming Organization (DSO). A well-established internal community accelerates adoption,

drives collaboration across teams, and transforms data streaming from a tactical tool into a strategic advantage.
By aligning IT teams, business units, and stakeholders around shared goals, the internal community can champion use cases, share knowledge, and position data streaming as the backbone of your enterprise. This alignment ensures that your organization moves faster, scales efficiently, and remains at the forefront of innovation.

Every journey begins with a first step

The time to act is now. Data streaming isn't just a technology—it's a paradigm shift that reshapes how businesses operate and compete. By starting your journey with a data streaming platform, you're not just adopting a tool; you're laying the foundation for a future-ready organization.

As the data streaming journey continues, remember that this is just the beginning. The landscape will keep evolving, and with it, new opportunities will arise. Embrace the change, stay informed, and use the power of data streaming to drive your organization into the future.

But whether you're just exploring or already implementing, this is your call to action: take the first step today and lead your industry into tomorrow.

Contact Us

Comments and thoughts on the book are encouraged. Or do you have a compelling data streaming example you would like featured in the next version of the book? Please reach out at **use-case-book@confluent.io**

Stay connected with Kai Waehner on
LinkedIn **kaiwaehner**
X **kaiwaehner**
for ongoing discussions and updates.

For regular insights into data streaming trends, architectures, use cases, and success stories, follow Kai Waehner's blog
www.kai-waehner.de

You can subscribe to his newsletter at
www.kai-waehner.de/news

For those interested in building a data streaming community and hear more about examples in this area, please reach out to Evi Schneider via
evi@confluent.io

The Ultimate Data Streaming Team

If you want to become a part of it, don't forget to reach out to use-case-book@confluent.io

Creating a book like The Ultimate Data Streaming Guide takes obviously a village and we are filled with immense gratitude for the collective effort that made this project possible. This work is the result of the collaboration, expertise, and encouragement of many individuals who contributed their time, knowledge, and support.

Thank you to our contributors, especially our customers, for sharing your stories and shaping this guide into a valuable resource. **To our reviewers and editors**, your commitment to precision has been invaluable. And to our readers and the data streaming community, your passion for innovation drives us. This book is for you and because of you.

Even though we will not be able to list everyone, we want to shout out some key team members:

Our managers **Kyle Gudgeon and Andrew Sellers**, for having our backs and making us stick to our commitments all. the. damn. time.

Our project managers **Ronja Antesberger and Lisa Kummerer** for holding us all accountable and delivering by our deadlines.

Queen of Design **Elena Hasenbeck**, your creative approach transferred so many visions into actual graphics.

Our linguistic heroes **Florian Stocker and Richard Ruf** as well as special MVP **Janine Van Poppelen**—all on a mission to written consistency!

Kate Sukhova and Emma Griffiths as well as the whole customer marketing team for collecting so many great real-life examples that are the heart and soul of this book.

Lyndon Hedderly for not only inventing the data streaming maturity curve, but also creating our most favourite graphic on business value of data streaming—thank you for your valuable support on this project.

Cody van Kampen and the product marketing team for their support in building out a table of contents that resonates!

A ton of proof readers, including family members and friends, that did not hold back with feedback as well as motivation—equally.

We're already working on the next edition of the book, featuring new use cases, real-life examples, and more community success stories. Want to be part of it? Email us at *use-case-book@confluent.io*

Kai Waehner & Evi Schneider

"This book is a fantastic resource for anyone looking to dive deep into the world of data streaming. It captures not only the technology but also the stories behind its adoption and the value it delivers."

Jay Kreps, *Inventor of Apache Kafka, Co-founder and CEO of Confluent*