



Confluent Use Cases

Gaetan Castelein, VP of Product Marketing

Josh Treichel, Partner Technical Lead

Agenda



Introduction



Contextual Event-Driven Applications



Universal Event Pipeline

- Security information & event management
- Middleware modernization
- Hybrid cloud
- Mainframe offload
- Streaming ETL

Adoption of Event Streaming

60%

Fortune 100 Companies
Using Apache Kafka

GoPro

stripe

airbnb



fitbit

box

WIKIPEDIA
the Free Encyclopedia

Pinterest

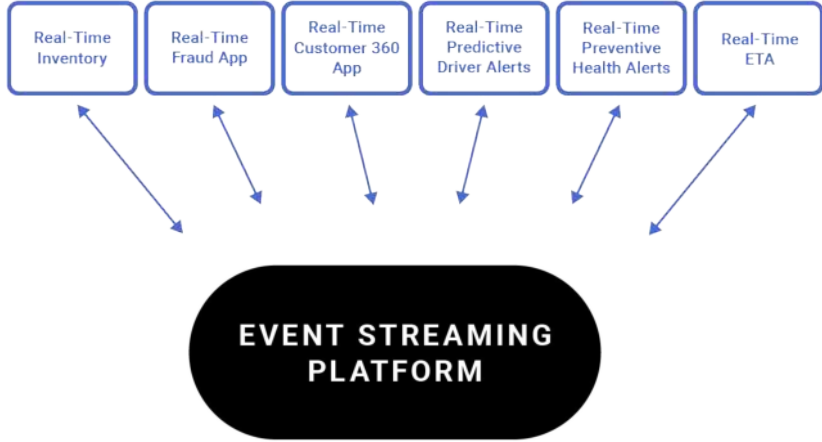
UBER

Event Streaming at the Heart of the Enterprise

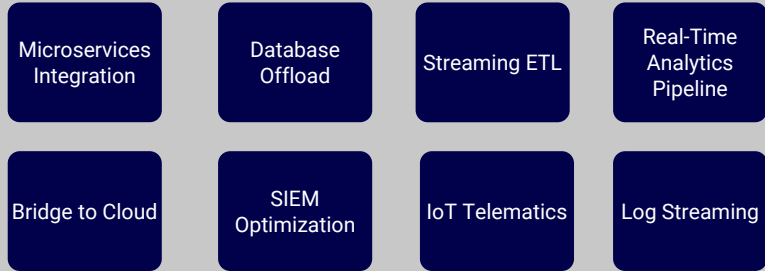


Contextual Event-Driven Apps

Industry verticals
App types



Pipeline use cases



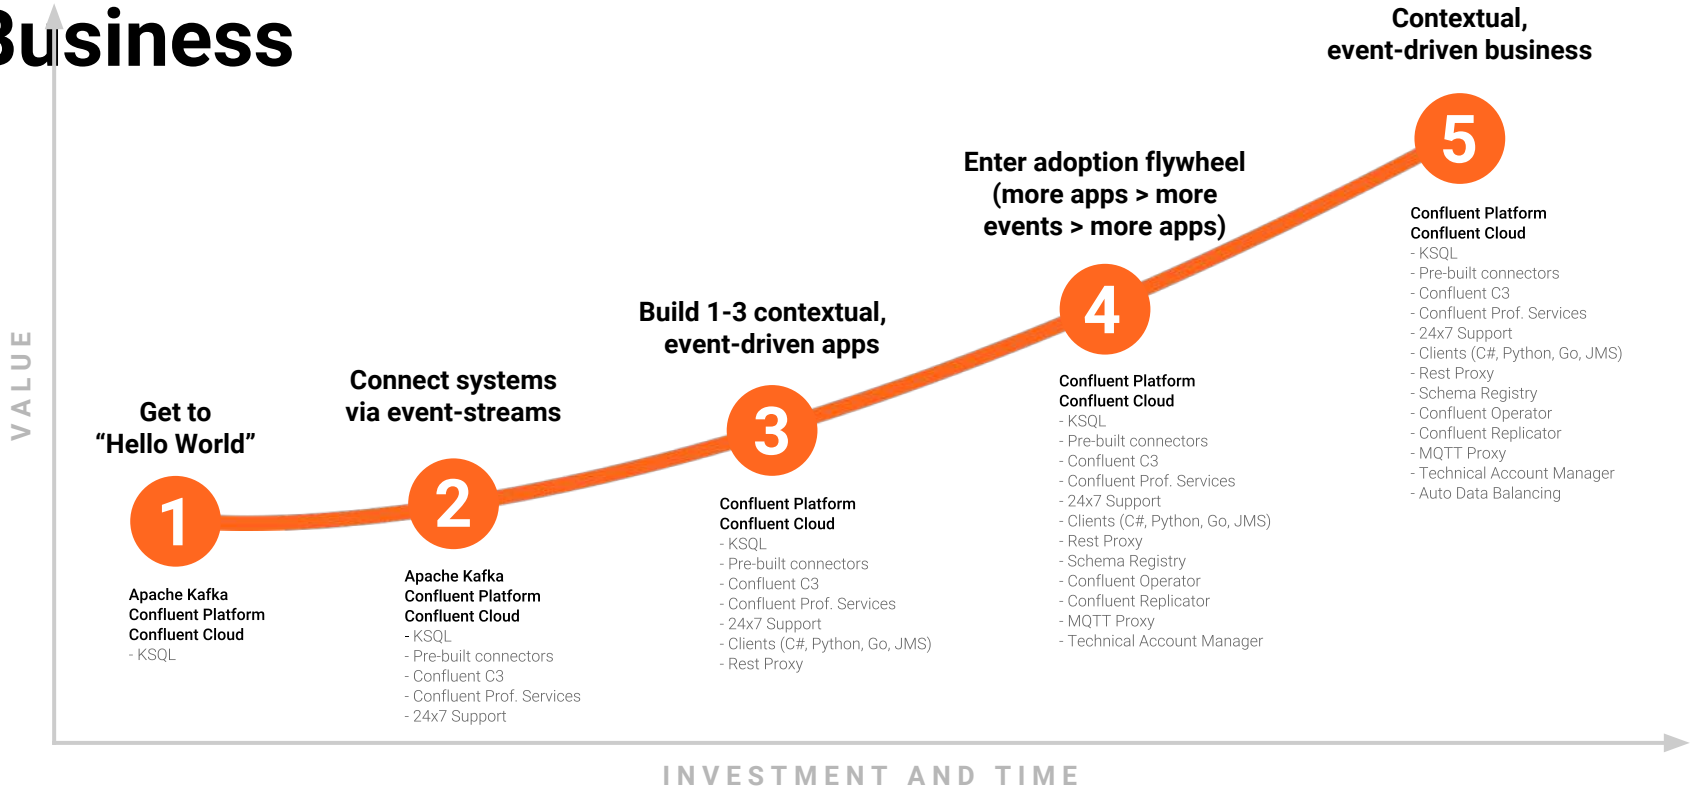
Paradigm shift

- Build new contextual event-driven apps
- Unlimited use cases, industry specific

Evolutionary

- Evolve the infrastructure
- Horizontal and repeatable use cases

Journey to the Contextual, Event-Driven Business



Agenda



Introduction



Contextual Event-Driven Applications



Universal Event Pipeline

- Security information & event management
- Middleware modernization
- Hybrid cloud
- Mainframe offload
- Streaming ETL

Capital One Customer 360



The screenshot shows the Capital One website interface. At the top, there is a navigation bar with the Capital One logo and links for Personal, Business, and Commercial. Below this is a secondary navigation bar with links for Credit Cards, Bank, Borrow, Invest, Learn, and Contact. The main content area features a breadcrumb trail: Capital One Home > Credit Cards > Benefits > Second Look. The title of the page is 'Second Look from Capital One'. Below the title is a large image of a man sitting in a chair, looking at a laptop. The text overlaid on the image reads: 'Helping you spot potential mistakes and unexpected charges'. Below the image, there is a sub-headline: 'Get a heads-up with Second Look® from Capital One®'. A paragraph of text follows: 'When Second Look identifies a charge on your account that looks like it might be a mistake—like a duplicate charge or an increase in a subscription charge—we'll send you an alert.* Should that charge turn out to be incorrect, we'll help you get it resolved.'

Customer problem

- $\frac{2}{3}$ of customers don't check their statements closely
- Duplicate charges, high tips, increased recurring charges


Prevent \$150 of fraud per year per customer

- The right level of signal vs noise for the consumer

Real time data cluster



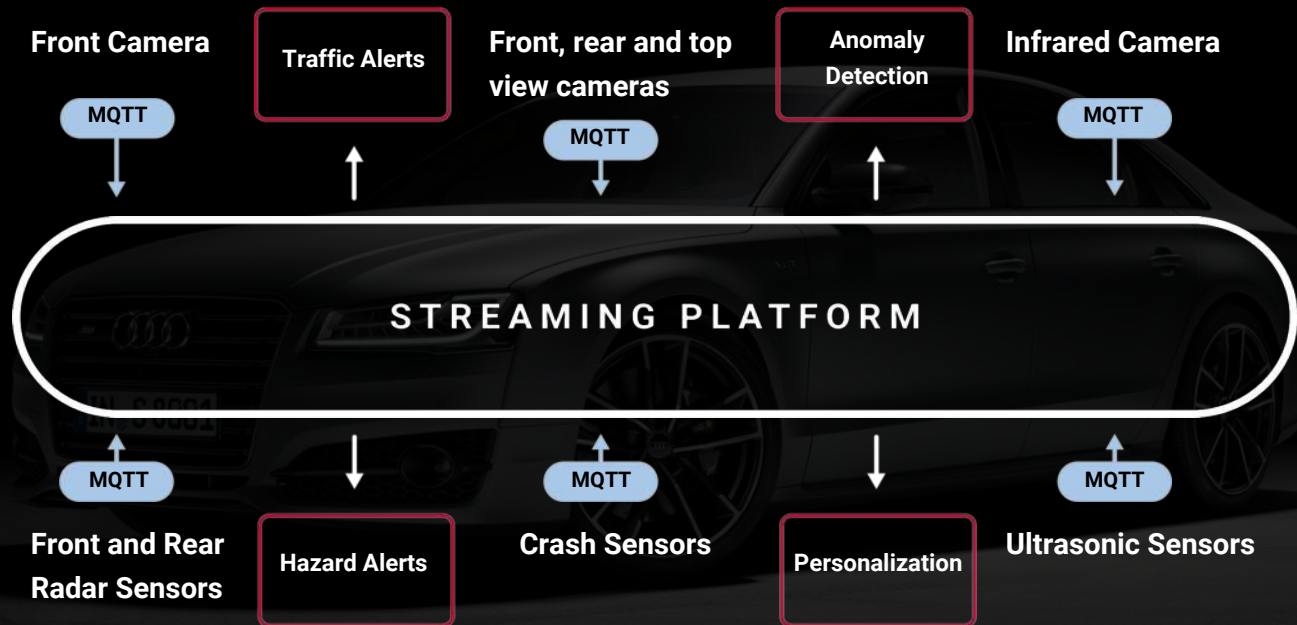
Audi




Real time data cluster



Audi



MiFID II compliance



Nordea

Kafka-powered MiFID II compliance

Nordea was able to reduce their platform costs by 73%, drop analytics turnaround time from 16 weeks to instantaneous reporting, and is now able to give all analysts access to trade data in real-time so they observe important patterns in data and respond to them in real-time.

Agenda



Introduction



Contextual Event-Driven Applications

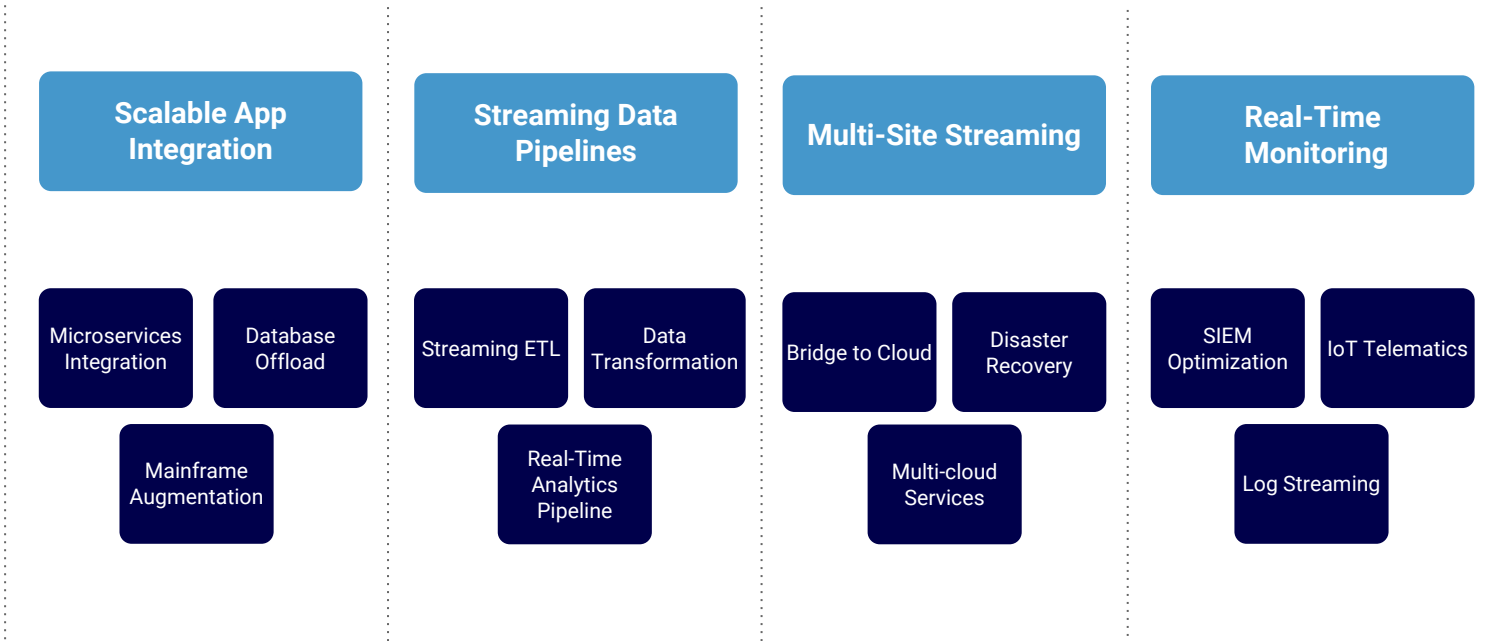


Universal Event Pipeline

- Security information & event management
- Middleware modernization
- Hybrid cloud
- Mainframe offload
- Streaming ETL

Solution Framework - Pipeline Use Cases

Universal Event Pipeline
Event-Enable the Enterprise



Pipeline Use Case #1

SIEM Optimization

SIEM Optimization

Security Information and Event Management

Help Cybersecurity teams expand their Security Operations Center investments to increase capability & reduce OpEx

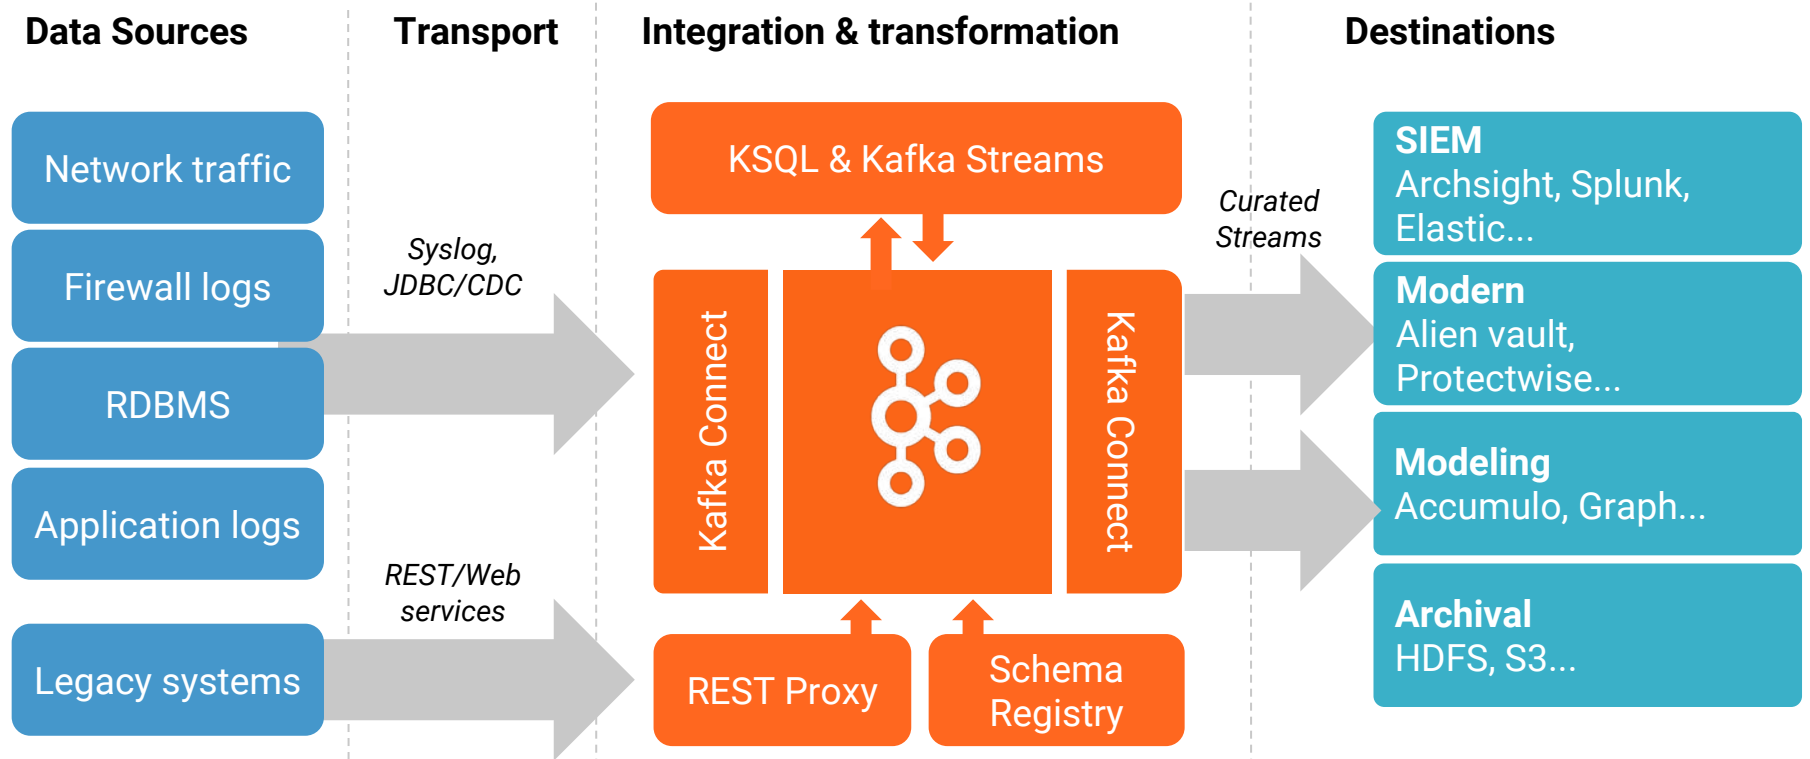
Customer Pain Points

- Cost, inflexibility of legacy SIEM
- Explosive data growth => \$\$\$
- Slow response to rapidly evolving threat matrix
- Difficulty expanding / augmenting Splunk or other SIEM

Business Value

- Fast reaction to new threats
- Migration to Elastic, others
- Infrastructure Consolidation
- Reduced OpEx

Sample SIEM Optimization architecture



Confluent solution:

- Confluent Kafka, KSQL and C3
- Integrated ArcSight, Elastic (ELK)

Application Use Cases:

- Modernized SIEM
- Opex reduction for Splunk & ArcSight

Infrastructure Use Cases:

- 100s of sources, 1000s of endpoints

Government Agency SIEM Modernization

Challenges:

- Huge cost with Arcsight, Splunk and Elastic (ELK)
- Vendor lock-in and capped data levels
- Capability limitations in any one SIEM tool
- Needed fast, real-time data movement
- Multi-site deployment across all combatant commands

Solution:

- Integrated ArcSight (legacy) and Elastic (ELK) using Kafka
- Integrated Kstreams and C3
- Real-time streaming analytics and processing of data

Pipeline Use Case #2

Microservices Integration

Microservices Integration

Modernize their legacy messaging solutions to a highly distributed, scalable, and real-time streaming platform

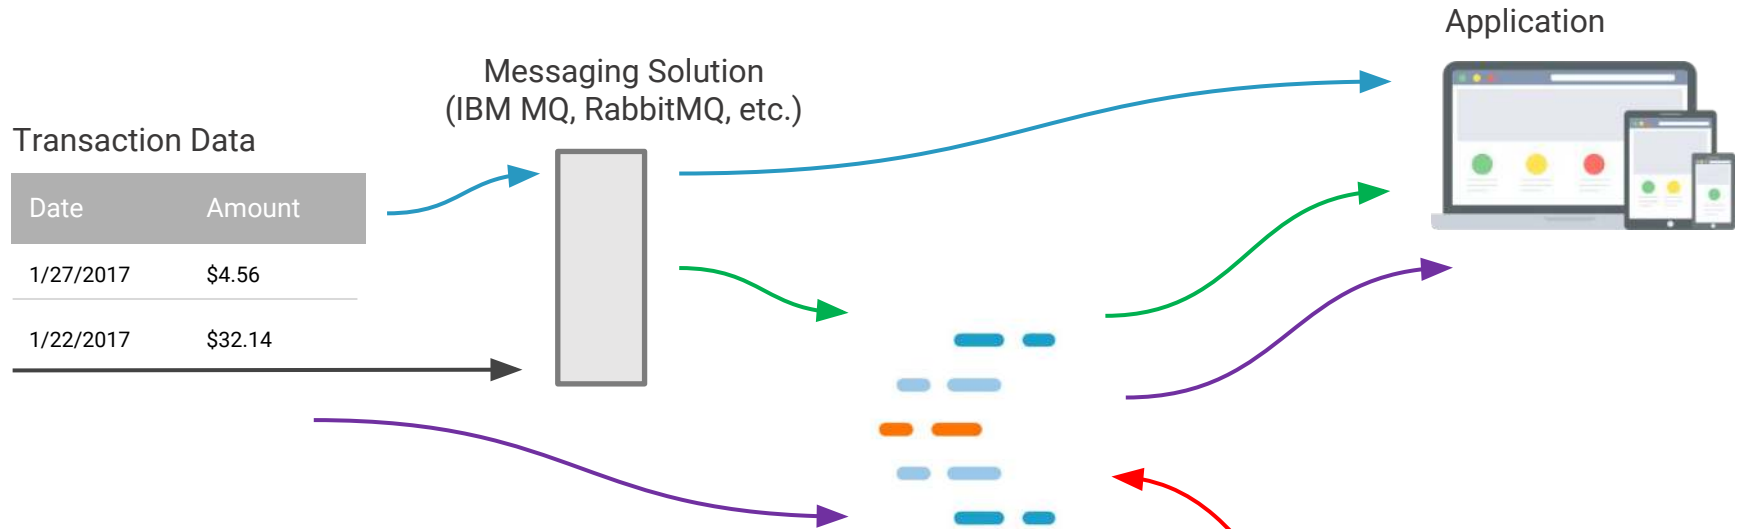
Customer Pain Points

- Scalability, low throughput
- Single point of failure
- Non persistent data
- Data integration complexity,
- Slow development velocity
- High cost of legacy systems such as mainframes

Business Value

- Reduce cost and complexity
- Accelerate development velocity by decoupling apps and data sources
- Offload messaging from legacy mainframes
- Build innovative real time apps

MQ Integration and Replacement



- 1) Direct Legacy MQ Communication with App
- 2) Kafka for decoupling between MQ and App
- 3) Direct communication via Kafka (no MQ anymore)
- 4) New projects and applications
(independent or related to the existing migration projects)

Kafka Microservices
Agile, small, lightweight (but scalable, robust)
Kafka microservice

External Solution
Big Data project (Elastic, Spark, AWS Sagemaker, ...)

Customer Stories



Moved off RabbitMQ to Confluent Platform to increase throughput and reduce mainframe costs. Confluent Platform manages 10s of thousands of requests per second to streamline the supply chain, manage inventory in real time, and support microservices

**Top 10
Global Bank**

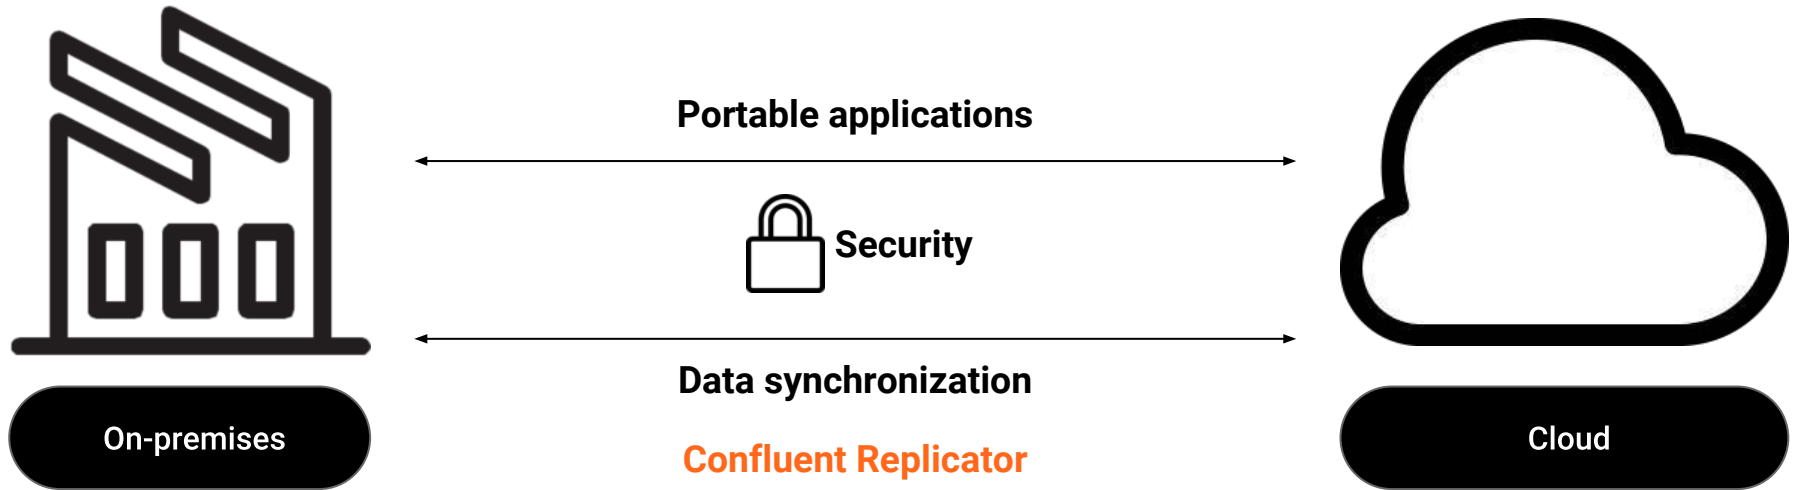
Moved off of IBM MQ and mainframes that were too costly and lacked real time streaming capabilities. Uses Confluent Platform to build a new ecommerce rewards platform that awards points for customer transactions in real time

Pipeline Use Case #3

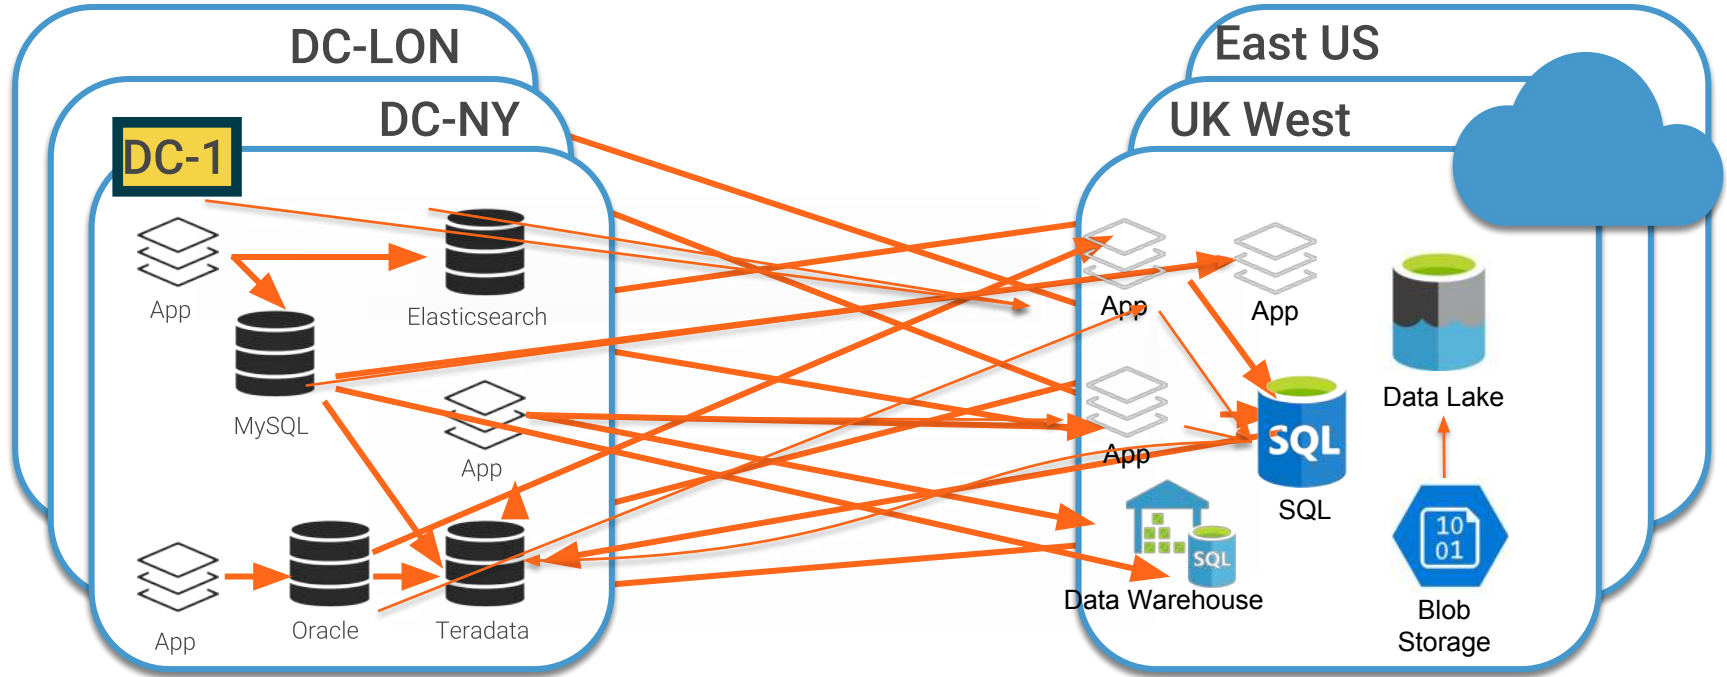
Bridge to Cloud

What is Hybrid Cloud?

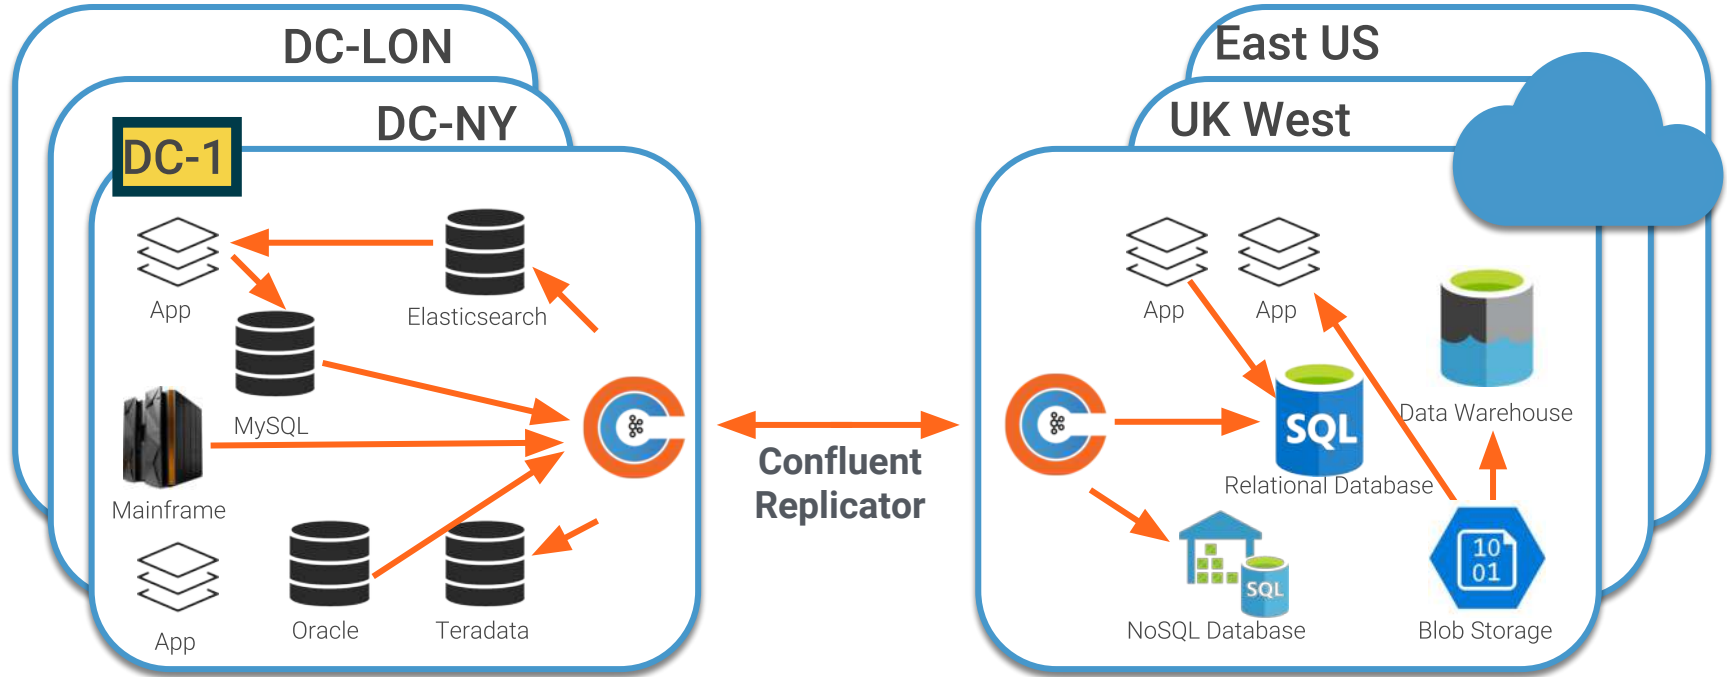
- **Market Definition** - mix of on-premises, and public cloud services with orchestration between the two



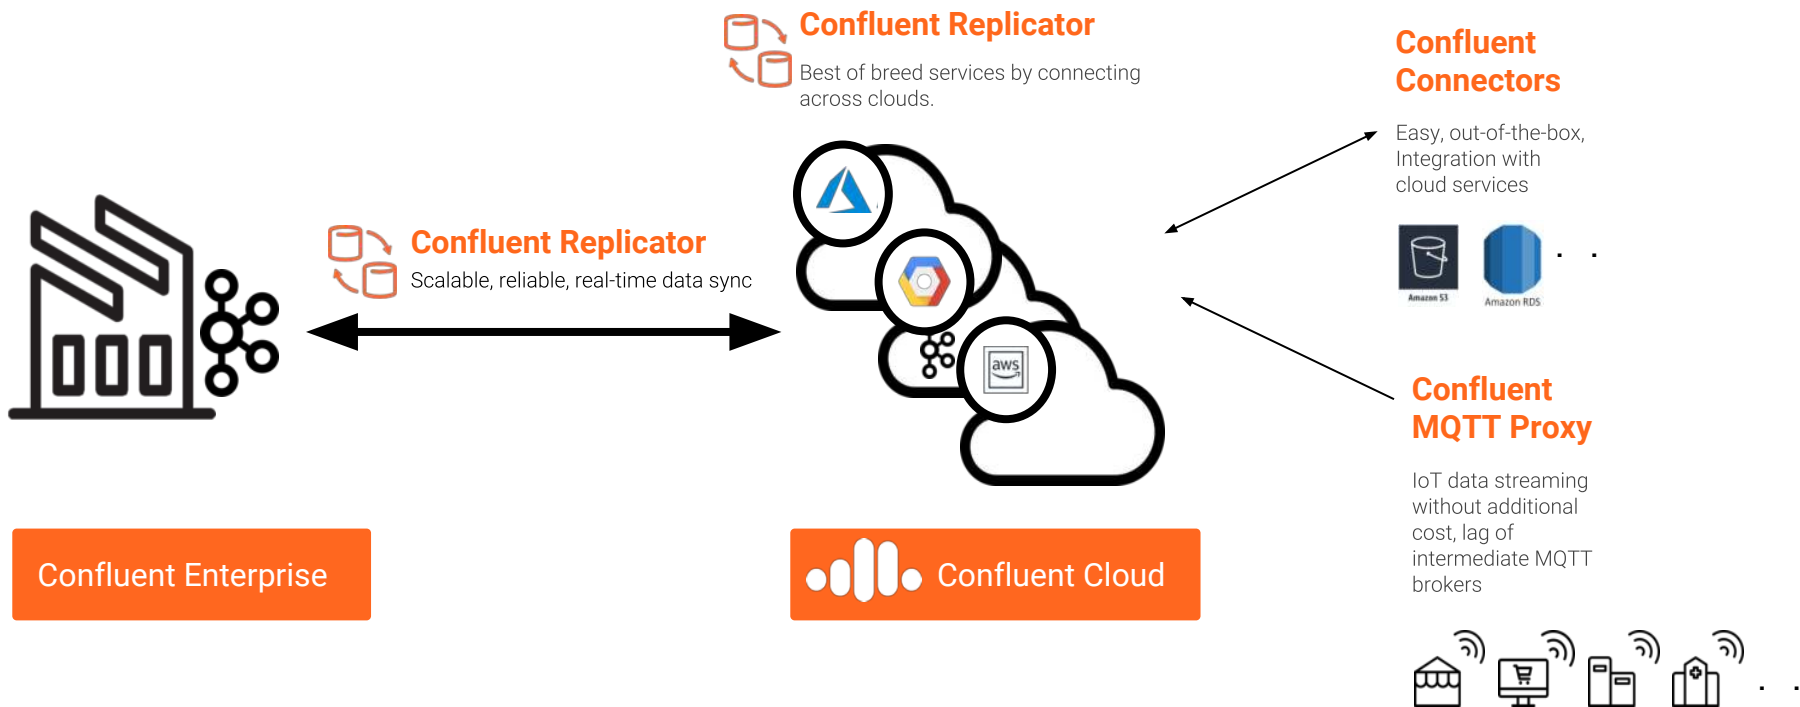
Hybrid cloud can get messy...



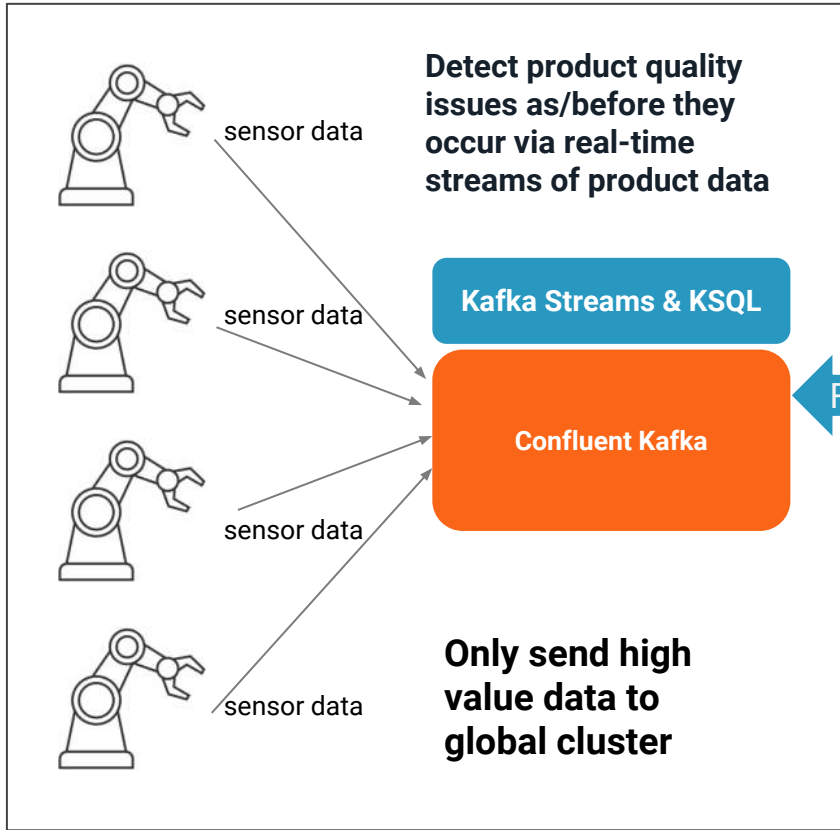
That's better



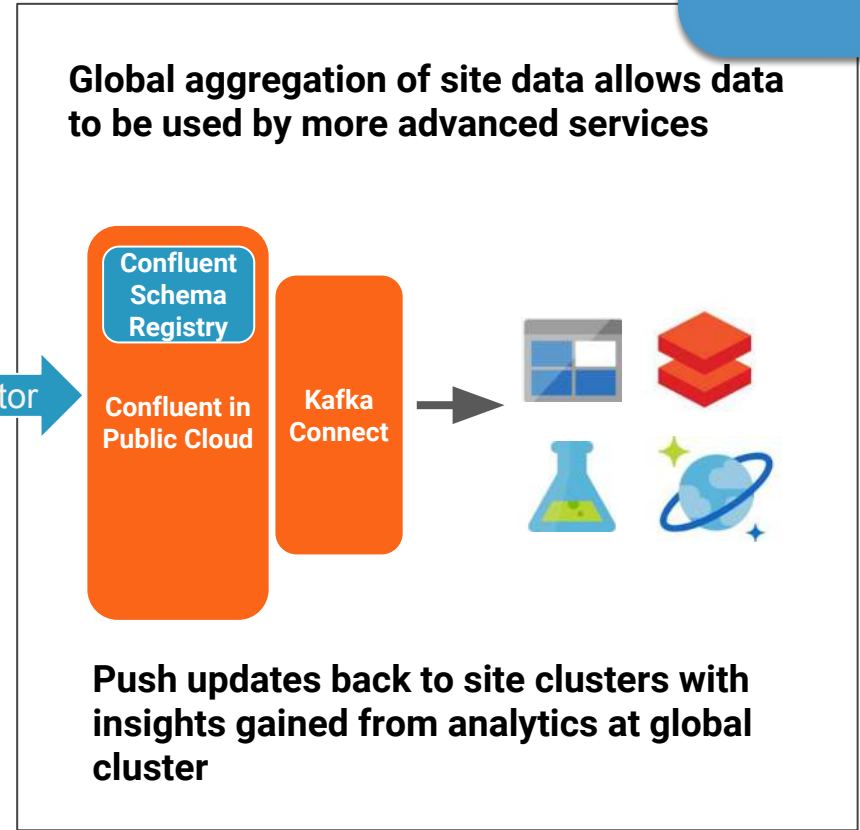
Industry's only hybrid Kafka service



Toyota Manufacturing Sites



Cloud Aggregation Cluster

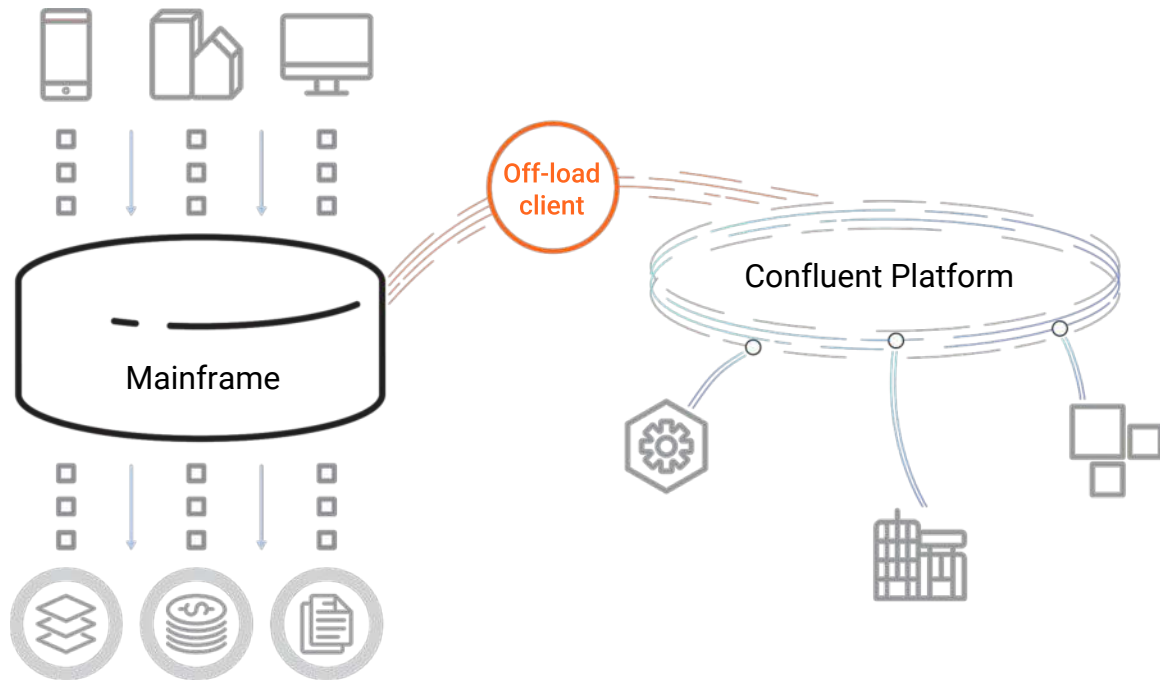


Pipeline Use Case #4

Mainframe Augmentation

Mainframe offload

Extend existing mainframe application with microservices by unlocking the data



Significant read-only transactions run through the Kafka Streaming Platform - reducing 30-40% of MIPS on mainframe

Top 10 Bank: Middleware modernization

\$25 million in savings through reduced mainframe MIPS


Transaction Data

Date	Amount
1/27/2017	\$4.56
1/22/2017	\$32.14

Transaction Description

Vendor	Description
Starbucks	Coffee
Walmart	Blu-Ray

Schema



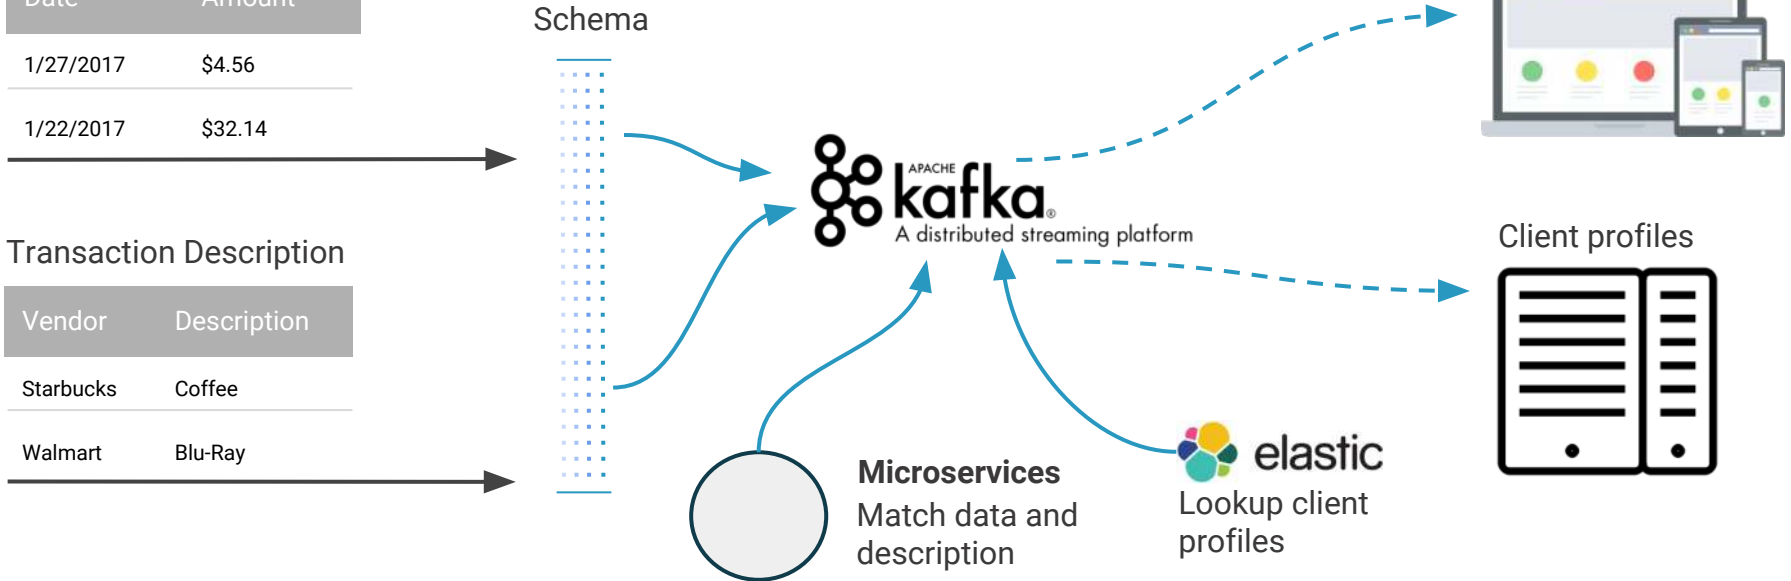
Microservices
Match data and description



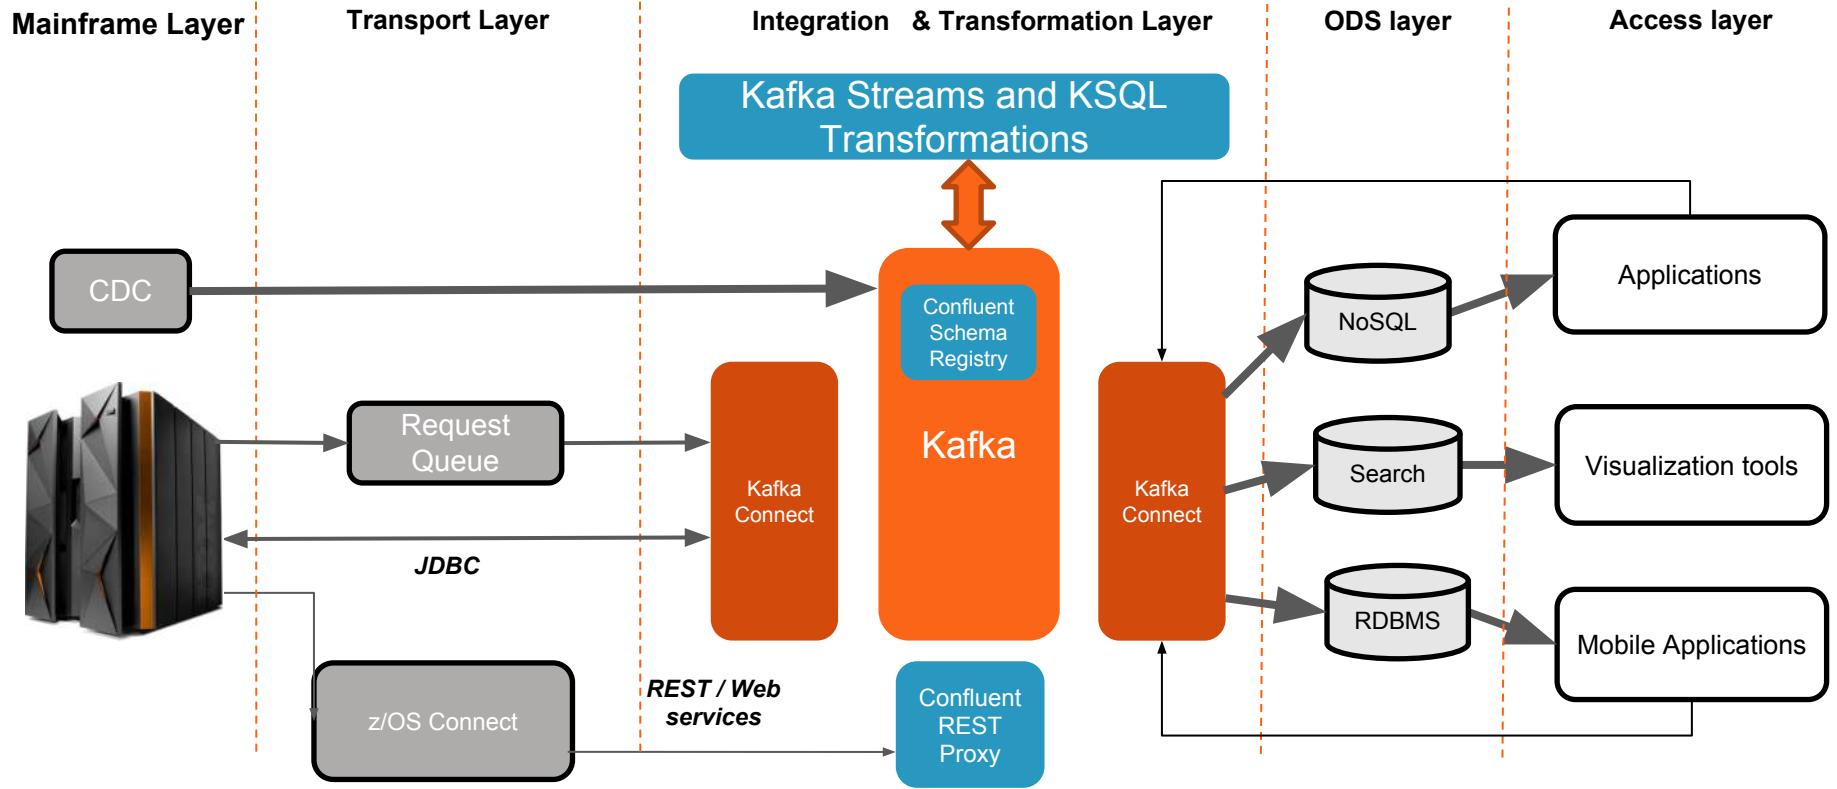
Website



Client profiles



Mainframe Read Offload - Sample Architecture



Pipeline Use Case #5

Streaming ETL

Streaming ETL

To exploit data, customers must:

- “Extract” it from transactional databases
- “Transform” it according to business / technical needs
- “Load” it into data warehouses/lakes for analysis

Customer Pain Points

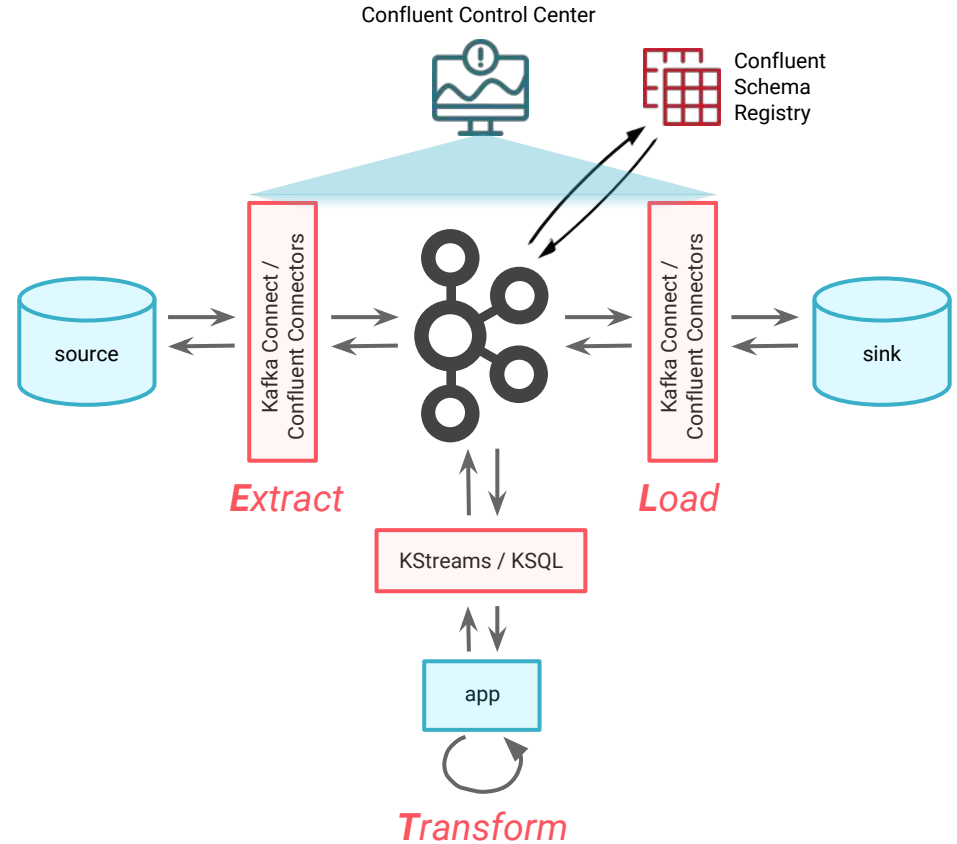
- High latency in reporting / analytics
- Brittle point-to-point connections
- Complex and time-consuming ETL processes
- Difficulty to change and scale data pipelines

Business Value

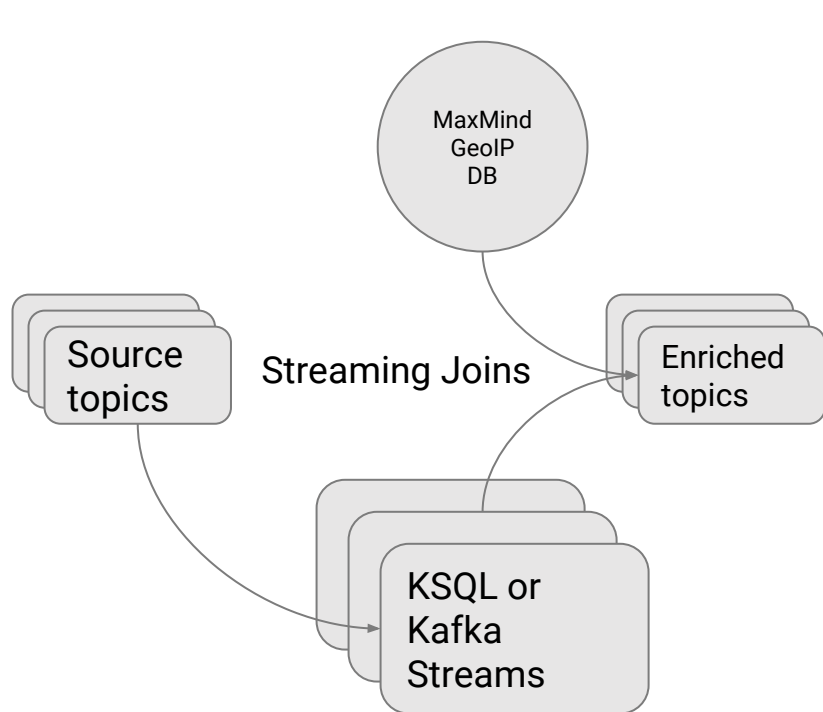
- Real-time reporting / analytics for accurate business decisions
- Lower licensing cost
- Streamlined processes
- Flexibility to change and grow without impact

Solution Overview

- Process **data in real-time**, or in batch (data persistence) where latency is not a concern
- Integrate all data sources and sinks in a **single platform** and avoid proliferation of data silos
- Decouple data producers from consumers to enable **maximum flexibility** without impact



Global Bank - Real-time GeoIP event enrichment



Kafka Streams app that reads from a (5) source topic(s) and appends the IP Geo location data and writes out to the final topic(s) that will be consumed by several enterprise fraud services and Enterprise Data Lake.

The logo for Confluent, featuring a stylized icon of horizontal lines of varying lengths on the left, followed by the word "confluent" in a lowercase, sans-serif font.

confluent