



***Migrate and Modernize Your Data  
Platform with Confluent and  
Databricks***



Data is changing the world



NETFLIX

Browse ▾


DVD

Search



Joshua ▾

### Top Picks for Joshua



### Trending Now




# Rich front-end experiences

### Because you watched Narcos



### New Releases





A wide-angle, high-angle photograph of a large industrial warehouse. The space is filled with a complex network of conveyor belts, primarily in shades of grey and yellow. The belts are arranged in multiple lanes, some curving and others straight, creating a sense of organized chaos. The floor is a light-colored, possibly concrete or metal grating. In the background, rows of tall shelving units are visible, densely packed with cardboard boxes. The lighting is bright and even, highlighting the scale and complexity of the facility. The overall atmosphere is one of a busy, modern manufacturing or distribution center.

Real-time backend operations

# Shared use cases across industries



## Retail

Customer 360  
Omnichannel optimization  
Supply chain optimization

**Improve customer attraction & conversion to increase revenue**



## Financial Services

Risk Management  
Trading Operations  
Security Operations & Fraud detection  
Payments Processing  
Customer 360

**Improved customer outcomes, increased revenue**



## Media & Telecom

Fraud Detection  
Customer churn reduction  
Real-time customer analytics  
Streaming health

**Better operational intelligence**



## Healthcare and Health sciences

IoT device management and analytics  
Clinical and claims

**Improved patient outcomes and cost effectiveness**



## Manufacturing / Automotive

Connected and autonomous cars  
Supply chain inventory  
Fleet management

**Predictive analytics generate exponential value**

# Companies are moving to the cloud to build modern data platforms



Agility



Elastic



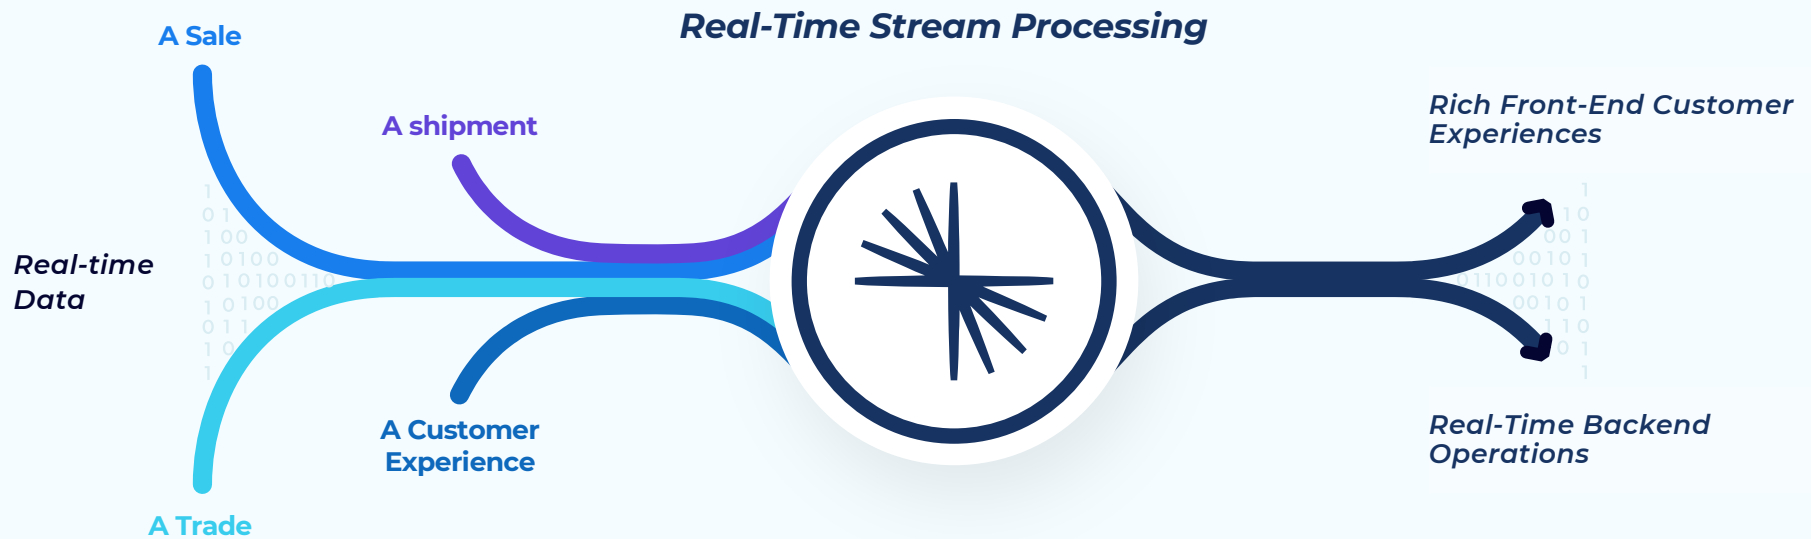
Cost-saving



Fast  
Deployment

*“More than 45% of IT spending on infrastructure and application software will shift from on-prem to cloud by 2024.” – **Gartner***

# Data streaming and Apache Kafka is required for a modern cloud data platform



**“We need to shift our thinking from everything at rest, to everything in motion.” — citi**

# Confluent Cloud

Help enterprises migrate and modernize their data platforms to power real-time analytics and applications



## Cloud-Native

Apache Kafka, fully managed and re-architected to harness the power of the cloud



## Complete

Go beyond Kafka to build real-time apps on the cloud quickly, reliably, and securely



## Everywhere

Maintain flexibility with deployments on any cloud and connect across clouds and on-premises that sync in real-time

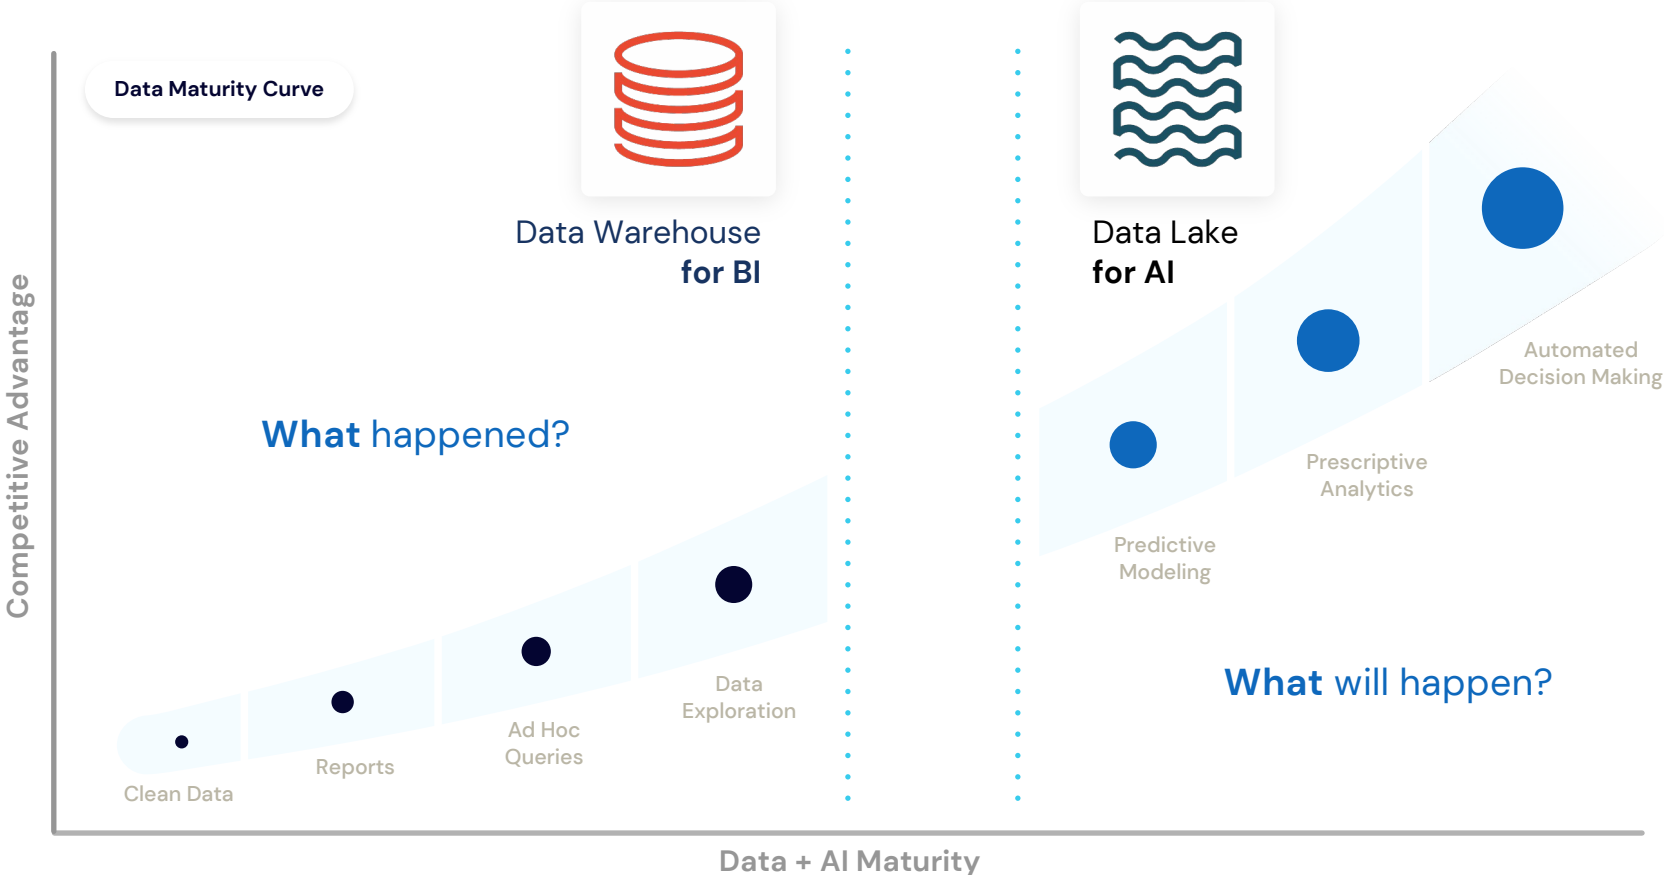
Stream confidently on a highly available platform built by the world's utmost Apache Kafka experts with built-in enterprise-ready security, compliance, and privacy controls



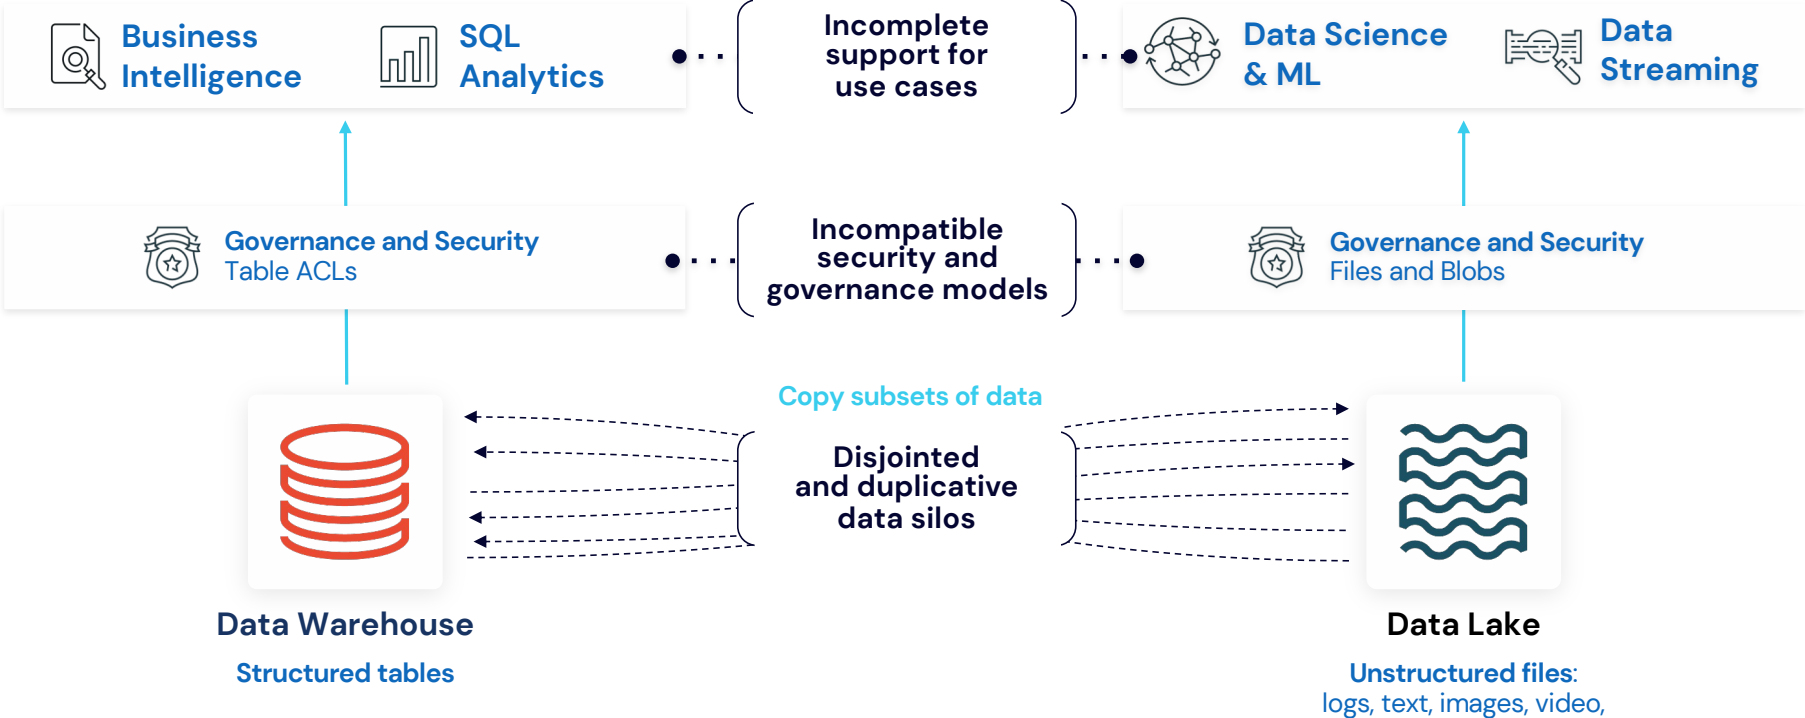


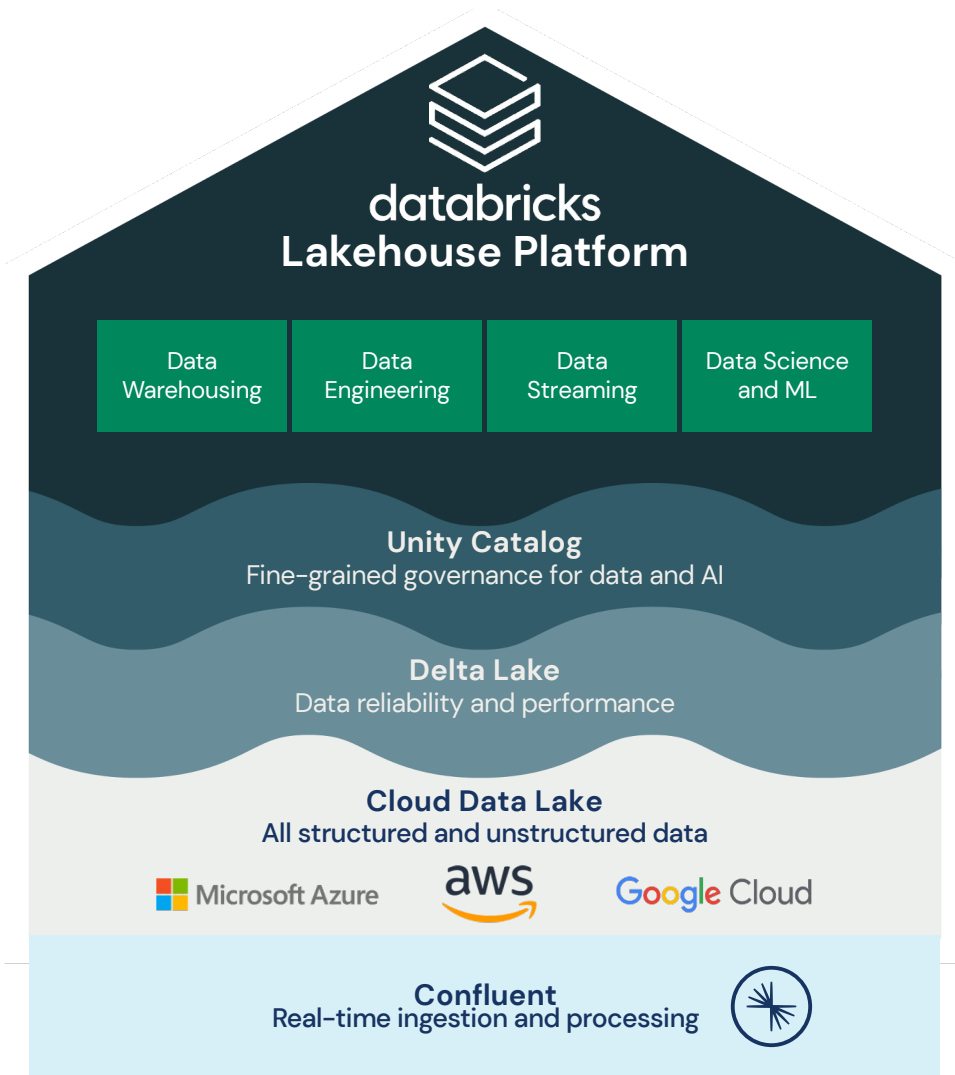
**Most enterprises still  
struggle with data, analytics,  
and AI**

# Realizing this requires two disparate, incompatible data platforms



# Realizing this requires two disparate, incompatible data platforms





# ***Databricks Lakehouse Platform***

## **Simple**

Unify your data warehousing and AI use cases on a single platform

## **Open**

Built on open source and open standards

## **Multicloud**

One consistent data platform across clouds



# Joint Value Proposition



**Confluent makes analytics on Databricks better by bringing in data in real-time**

A complete **real-time** event streaming and **analytics** solution for all data, analytics and AI workloads



BI and reporting in the lakehouse have more recent and valuable data for better decision making



ML/AI models can get real-time data from apps, IoT and consumer interactions for better predictions

*Delivered at scale with the speed, security, and reliability required by enterprises*

***From  
migration to  
modernization***

How Confluent and Databricks can  
accelerate your journey to real-time  
analytics in the cloud



# Migrating and modernizing your data warehouse is not easy

## ✗ High Cost

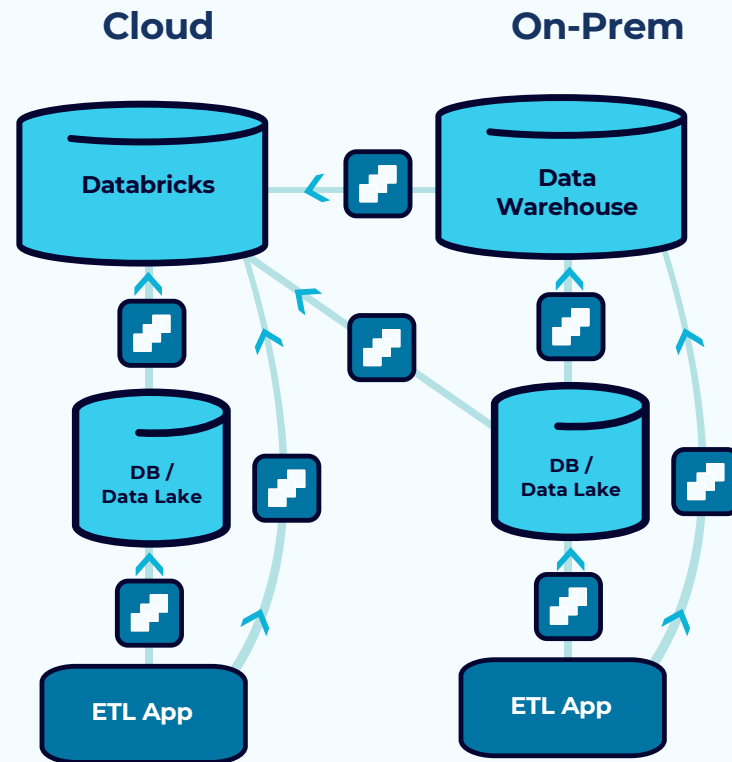
More data systems (e.g. DB) or DW costs to handle volume of data.

## ✗ Slow time to analytics-ready data

Batch based processes with limited scale for “real-time”. Migrating DW to cloud can take years.

## ✗ No support for hybrid/multicloud

Having distributed data and systems across environments compounds issues.



 = batch jobs and APIs/ data copying

# Modernize your data warehouse with Confluent and Databricks



## Reduce total cost of ownership

**Process data in stream** to lower DW costs. Lower data pipeline TCO with fully managed cloud service.



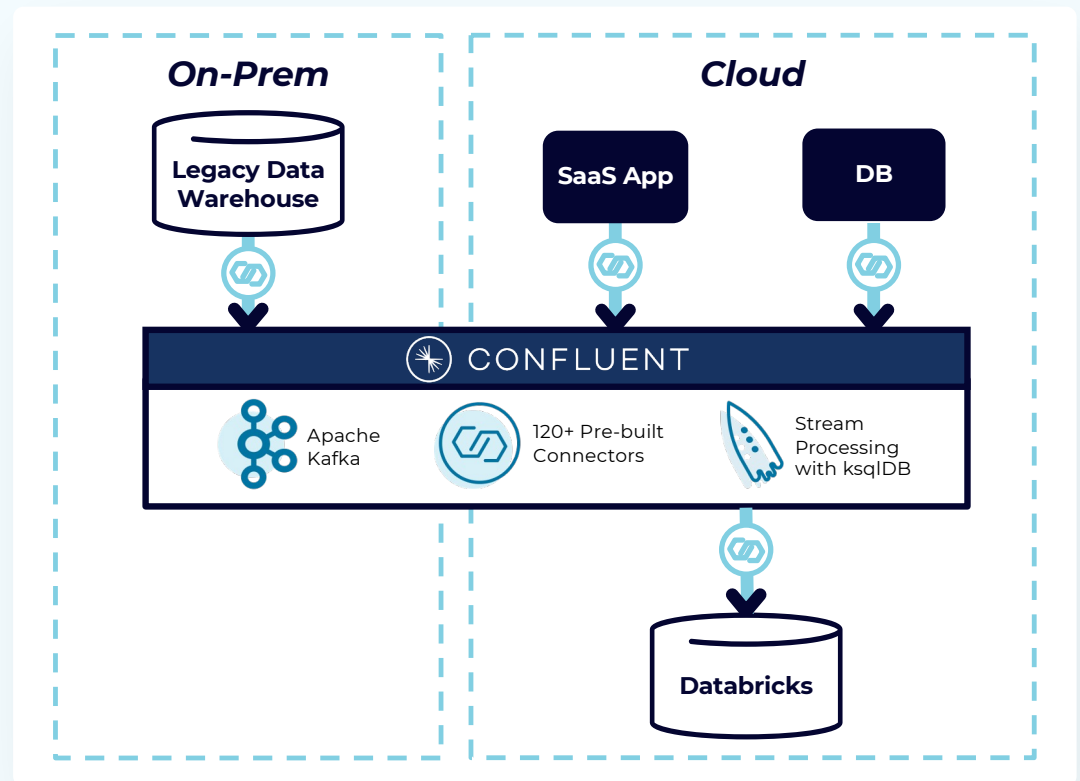
## Power new analytics and apps

Link on-prem and cloud for easier data movement across environments with **real-time event streaming** to modernize your analytics and applications



## Get more data to and from your DW

**Break data silos** and easily connect your data warehouse to popular sources and sinks using Confluent's 120+ pre-built connectors



Real-time connections & streams



## Sources



IoT



Legacy Data Stores:  
Netezza, Teradata  
Oracle, Mainframes  
Databases

## Event Streaming and Processing

Kafka Streams  
& ksqlDB - real-time  
stream processing  
and transformations

Event Streaming  
Platform  
built on Kafka



## Analytics



Databricks Data  
Science Workspace

TensorFlow PyTorch

Databricks BI  
Workspace

Looker

Data Lake

**DELTA LAKE**  
Delivers reliability, performance  
and governance on your data lake

Data Warehouse



On Premises  
or any cloud

On Premises  
or any cloud

On any  
cloud

# Instantly Connect Popular Data Sources & Sinks

**120+**  
pre-built connectors

**databricks**

Logos shown include: amazon REDSHIFT, amazon KINESIS, amazon S3, Microsoft Azure Blob Storage, ATTUNITY, camunda, Couchbase, Azure Data Lake, cassandra, Apache ACTIVEMQ, elastic, DATASTAX, GridGain, HumIO, AWS Lambda, Google Cloud Storage, HADOOP, APACHE HBASE, HVR, IBM, imply, HTTP, IBM MQ, JDBC, JMS, kinetica, XENON, mongoDB, MAPR, Microsoft SQL Server, MQTT, omni-sci, neo4j, ORACLE GOLDEN+GATE, Pivotal Greenplum, PUSH TECHNOLOGY, PostgreSQL, RabbitMQ, salesforce, splunk, syslog, [ROCKET], SAP HANA, snowflake, VERTICA.

100+ Confluent Supported

30+ Partner Supported, Confluent Verified

# Confluent: Everywhere

Fully-Managed Service

## Confluent Cloud

Apache Kafka Re-engineered  
for the Cloud



Available on the leading public clouds



Self-Managed Software

## Confluent Platform

The Enterprise Distribution of  
Apache Kafka



Deploy on any platform, on-prem or cloud



# 3 Modalities of Stream Processing with Confluent



## Kafka clients

```
ConsumerRecords<String, String> records = consumer.poll(100);
Map<String, Integer> counts = new DefaultMap<String,
Integer>();
for (ConsumerRecord<String, Integer> record : records) {
    String key = record.key();
    int c = counts.get(key)
    c += record.value()
    counts.put(key, c)
}
for (Map.Entry<String, Integer> entry : counts.entrySet()) {
    int stateCount;
    int attempts;
    while (attempts++ < MAX_RETRIES) {
        try {
            stateCount = stateStore.getValue(entry.getKey())
            stateStore.setValue(entry.getKey(), entry.getValue() +
stateCount)
            break;
        } catch (StateStoreException e) {
            RetryUtils.backoff(attempts);
        }
    }
}
```

## Kafka Streams

```
builder
    .stream("input-stream",
        Consumed.with(Serdes.String(), Serdes.String()))
    .groupBy((key, value) -> value)
    .count()
    .toStream()
    .to("counts", Produced.with(Serdes.String(), Serdes.Long()));
```

## ksqlDB

```
SELECT x, count(*) FROM stream GROUP BY x EMIT CHANGES;
```



Flexibility

Simplicity



# Deliver value faster with *ksqlDB Recipes*

Delve into our collection of top stream processing use cases, which provide code samples and step-by-step tutorials to help you from idea to proof-of-concept.

Plus, to make it even easier, there is a one-click button that will launch the SQL script directly within Confluent Cloud.

[developer.confluent.io/tutorials](https://developer.confluent.io/tutorials)



## Anomaly and Pattern Detection

Identify patterns and unusual activity to detect fraud and monitor network health

[Detect Unusual Credit Card Activity](#)

[Flag Abnormal IoT Devices](#)

[Handle Corrupted Data from Salesforce](#)



## Customer 360

Combine data sources for a unified 360-degree view of the customer across channels

[Match Mutually Interested Online Dating Users](#)

[Understand User Behavior with Clickstream Data](#)

[Notify Passengers of Flight Updates](#)

[Build Customer Loyalty Schemes](#)

[Run analysis of a customer's online journey](#)



## Cybersecurity

Protect critical systems and sensitive information by detecting threats in real-time

[Detect and Analyze SSH Attacks](#)

[Detect a slow DDoS attack](#)

[Identify firewall deny events](#)

[Filter Audit Logs for Output to Splunk](#)



## Logistics and IoT Management

Optimize supply chain operations and detect future failures with predictive maintenance

[Optimize Fleet Management](#)

[Track order shipments in real-time](#)



## Predictive Analytics

Train predictive models by using machine learning to anticipate future outcomes

[Optimize Omni-channel Inventory](#)

[Create personalized banking promotions](#)

[Retrain a machine learning model](#)



## Real-Time Alerts and Notifications

Build real-time alerts and notifications to power digital customer experiences

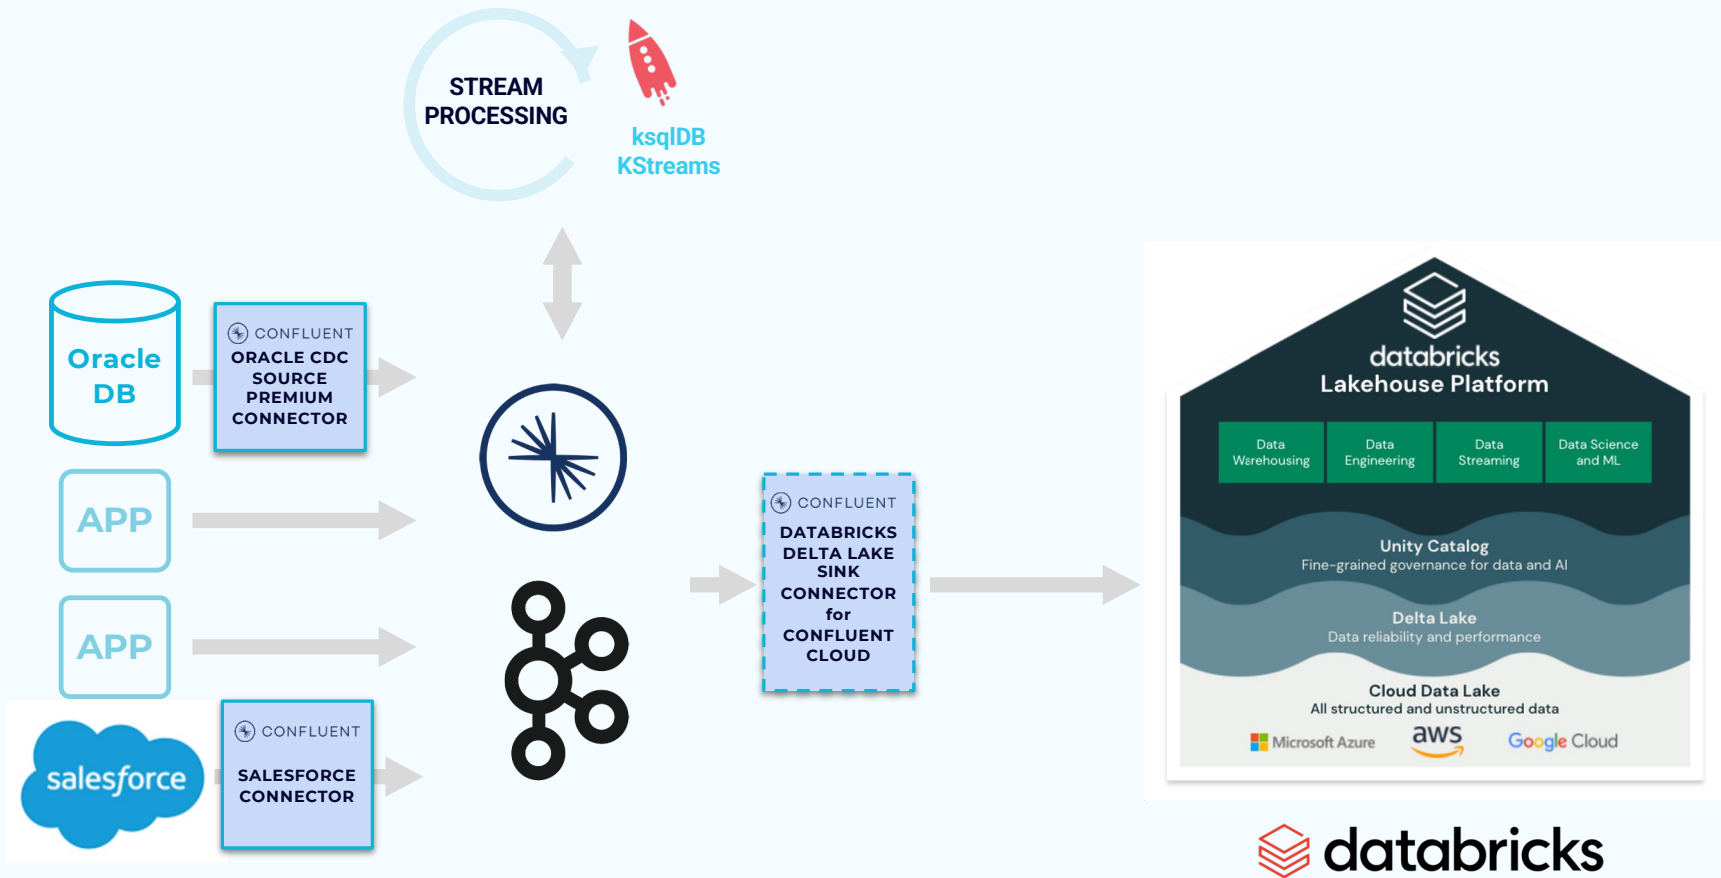
[Create geolocation alerts and promotions](#)

[Notify Passengers of a Flight Delay](#)

## *Demo*

A financial service company wants to build real-time dashboards and analytics in Databricks but their data is spread across an on-prem Oracle data warehouse and cloud system.





**Q&A**

Please drop your questions in the chat.





# Get started for **FREE**

**\$400 free** over  
**2 months** from sign-up

[confluent.io/confluent-cloud](https://confluent.io/confluent-cloud)



The screenshot shows the 'Developer' section of the Confluent website. At the top, there's a navigation bar with 'Developer' and a 'QUICK START' button. Below that, the main heading is 'APACHE kafka with Confluent' with the subtext 'The easiest way to learn Apache Kafka® and Event Streaming'. A 'QUICK START' button is visible. The main content area features a video player titled 'What is Apache Kafka?' by 'KAFKA CO-CREATOR JAY KREPS, CEO OF CONFLUENT'. The video thumbnail shows a man in a white shirt and the text 'WELCOME TO Confluent Developer CONFLUENT'. There are also 'Watch on YouTube' and 'Copy link' options.

[developer.confluent.io](https://developer.confluent.io)

The new one-stop shop for learning Apache Kafka, event streaming, and associated cloud technologies. The site comes with an extensive library of Kafka-related courses.



***Thank you!***

