

ROYAL BANK OF CANADA

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



Headquarters

Montreal, Canada

Industry

Financial Services

Challenge

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

Solution

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

Results

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



Confluent
Control Center



Multi-Datcenter
Replication



Schema
Registry

As Canada's biggest bank, and one of the largest in the world based on market capitalization, RBC has a diversified business model with a focus on innovation and providing exceptional experiences to their 16 million clients in Canada, the U.S. and 35 other countries. To continue the bank's mission of providing world-class products and services to its clients, RBC's Data and Analytics (DNA) and Enterprise Cloud teams joined forces to build a real-time, scalable and event-driven data architecture for the bank's growing number of cloud, machine learning and AI initiatives.

The DNA and Cloud teams partner with several lines of business across the bank such as wealth management, digital marketing, corporate real estate and fraud, by providing access to technology, knowledge sharing and training opportunities across the business groups. This partnership allows RBC to design new services faster and meet the needs of RBC's diverse client base. They also help the business lines bring up new services faster while helping the development teams navigate complex data ownership, security, cloud and regulatory compliance.

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

—
Mike Krolnick, Head of Engineering, Enterprise Cloud, RBC

To deliver the event-based architecture for RBC's business lines, the DNA and Cloud team selected Confluent Platform and Apache Kafka® to be the foundation of their scalable, real-time streaming data architecture. RBC now has a cross-functional team helping transform the bank to a technology and data-driven organization.

RBC's goal through this transformation is to capitalize on the next generation of open-source, cloud-enabled, real-time and responsive application and software development. By adopting Confluent Platform and Apache Kafka, RBC is realizing this objective, as is evidenced by their deployment of over 50 applications on their event-driven data architecture.

Due to regulation, security and the need for auditing, data within these various groups must be handled thoughtfully, with insight on how to maximize its value and quality without introducing risk.

CHALLENGES

Like many banks with a long history and a large client base, RBC's accumulated assets are very complex. As the business grows, technologies evolve and material events like acquisitions occur, RBC exerts force to adapt itself to these changes. RBC's data landscape exists across lines of business and within a centrally managed infrastructure. Due to regulation, security and the need for auditing, data within these various groups must be handled thoughtfully, with insight on how to maximize its value and quality without introducing risk to clients or the organization.

"We realize that to prepare for a more competitive future and to extend our leadership in innovation, we need to be able to generate accurate data and business insights," said Kerry Joel, Sr. Director, Product Innovation, Data and Analytics.

Part of the central technology stack at RBC is a mainframe-based infrastructure that supports many applications. This infrastructure is the system of record for a number of client profiles and before Confluent, accessing the data could be delayed, complex and costly for the bank.

"We needed a way to rescue data off of these accumulated assets, including the mainframe, in a cloud native, microservice-based fashion," said Mike Krolnik, Head of Engineering within the Enterprise Cloud team.

The mainframe serves data as clients search for information on their website. This produced a workload that was high in reads, low in writes. Also, because it processed new transactions using a batch-oriented database, providing real-time updates on account balances was limited. While maintaining the writes to the client profiles was necessary, the high read workload was a candidate for elimination.

SOLUTION

RBC selected Confluent Platform based on Apache Kafka, and wrote a microservice to significantly reduce the reads on the mainframe, saving RBC fixed infrastructure costs (OPEX). RBC stayed compliant with bank regulations and business logic, and is now able to create new applications using the same event-based architecture.

After a significant number of additional successful application deployments, RBC's DNA and Cloud teams implemented Confluent Platform to act as a centralized real-time data platform to support the bank's business lines. Confluent Platform flexibly ingests data from RBC's core banking platforms, such as the mainframe and data from their cloud-native applications. Ultimately, the bank has adopted an event-based architecture which is allowing developers, operators and the C-suite at RBC to re-imagine how data is stored and analyzed.

Kafka transforms even the most basic of initiatives.
Adoption of Kafka at RBC has been massive and organic.
Within the first six weeks after our launch of Kafka, we
had 37 teams asking to use Kafka for various projects
and initiatives.

—
Kerry Joel, Sr Director, Product Innovation, Data and Analytics

TECHNOLOGY PATTERNS REALIZED FOR OPTIMAL SUCCESS

RBC experienced the emergence of four technology patterns throughout this transformation:

1. The first was moving the accumulated assets to a cloud-native, microservices architecture without rewriting the core systems of record. This meant RBC could slice large monolithic applications into smaller, more agile components with real-time data synchronization. "We could begin the process of slicing our large, monolithic applications into smaller, more agile pieces," said Kroinik.
2. Using microservices, RBC built new functionality with event processors rather than a database or ETL patterns. This allows for more flexible and decoupled systems without needing to rewrite the entire system.
3. Business teams were now able to perform quick data discovery and analysis for data-driven insights. These insights span transactions, decision support and interactions in real time.
4. RBC can decouple everything to increase the speed of innovation and create new flows without impacting current operational systems, which is empowered by Confluent Platform's ability to orchestrate complex business flows.

The simplicity of Confluent added to RBC's decision to adopt the technology. "Our journey began when we looked into event-based architecture," said Kolnik. "We realized the potential of that architecture to transform our business."

"The ability to stream events transforms even the most basic of initiatives," said Joel. "Adoption at RBC has been massive and organic. Within the first six weeks after our launch of the Confluent Platform, we had 37 teams identified as early Kafka adopters for various projects and initiatives."

In addition, RBC relies on Control Center, a monitoring tool in Confluent Enterprise, that gives the various operating teams the ability and confidence to manage their environment company wide. "Control Center gives you full visibility of all messages as they move through the system, such as what the consumers and the producers are doing and whether they're there or not," said Krolnik. "This is the number one question we get as a platform support team."

To increase the speed of innovation, establishing the schema became a foundational part of RBC's success with Confluent. The ability to evolve the schema has allowed RBC to speed operations without impacting other consumers and producers. "Right now, most of the things we build are dependent on each other. If one changes the other needs to change. That slows down time to market, that slows down change, it slows down innovation," said Krolnik. "With Kafka and our Schema evolution strategy, we're able to decouple everything. So you can change either side of the equation. I think that's something unique in the platform."

The simplicity of implementing Kafka added to RBC's decision to partner with Confluent. "Most of the open source tools are challenging, but not Kafka," said Krolnik.

Confluent created an open source event streaming platform and reimagined it as an enterprise solution. We understand the capabilities it gives our business and implemented it within our enterprise data strategy.

—
Kerry Joel, Sr Director, Product Innovation, Data and Analytics

RESULTS

“Confluent created an open source event streaming platform and reimagined it as an enterprise solution,” said Joel. “We understand the capabilities it gives our business and implemented it within our enterprise data strategy.” The results RBC has experienced since implementing the streaming platform are:

- **Increased efficiency in app building.** RBC has built over thousands of apps across the business lines and has hundreds of projects running at any time across the bank. “Kafka is so much more efficient and stronger than what we were using before, and Confluent brings a lot more value to the way we build apps,” said Krolnik. “Kafka was different than I thought it was. Streaming eliminated the problems we had.”
- **Lowered anomaly detection time from weeks to real-time.** Before Kafka, RBC would spend weeks identifying data anomalies. “At RBC, accuracy of data is incredibly important. We spend a substantial amount of time asking ‘why is this data wrong?’ To have Confluent answer that question for us by using machine learning techniques to detect variances, we save many, many weeks of manual data analysis. That’s time earned back that can be spent on deriving business insights from that same data, instead of questioning it,” said Krolnik.
- **Implemented data reuse across teams for relevant business insights.** Immediately after developing RBC’s first Kafka use case, a successful sales enablement tool, the RBC team realized the power of true information sharing across the organization. “When we first put the data into Kafka we realized it was the first time we could really share data across teams. Kafka has created a lot of firsts for RBC,” said Krolnik. “With Kafka, we can have many people talking to the same data source without interfering with one another, which is a huge benefit for us.”

“RBC is a people business. As a financial institution, it’s always about our clients’ trust and safeguarding their information and money is always a top priority,” said Joel. “To manage that well, we focus on developing products and services that reflect how clients want us to interact with and support them.”

RBC is a people business. As a financial institution, it’s always about our customers and we’re humbled by the fact that we are entrusted to hold their money. To manage that well, we focus on developing products and services that reflect how customers want us to interact and support them.

–

Kerry Joel, Sr Director, Product Innovation, Data and Analytics