White Paper

Modernizing Financial Services Data Layer With Redis Enterprise

redis

Table of Contents

Introduction

As COVID and technological innovations have propelled businesses forward into a new digital era, customers now demand a digital banking experience that's fast, easy, and tailored to their needs. Banks are making the transition to provide a more digitized service but will need to revamp everything from strategy to approach through to the IT systems they use to meet the sky-high expectations of the modern consumer.

In the rush to deliver digital services and products, banks have made the modernization of front-end IT a priority. Based on a 451 Research, a part of S&P Global Market Intelligence Voice of the Enterprise: Customer Experience 2021 survey, the finance industry comes out as the top vertical to have a digital transformation strategy. In the study, 62% of surveyors claimed to be actively digitizing business processes and technologies and 61% claimed that improving the customer experience was the main driver for digital transformation.[1]

As the research indicates, maximizing the user experience is one of the greatest motivations behind digital transformation. But if that's the case, then why are banks failing to provide a digital banking experience that meets the demands of the modern consumer?

According to the CapGemini report, the failure to meet user expectations likely stems from neglecting the modernization of core middle and back-office transaction and operational systems. Many banks still follow outdated business processes that are incompatible in today's digital era. These processes still use rigid legacy systems, tedious paper-based documentation, and cumbersome manual processes that are inherently complex and siloed.

Half of banking customers surveyed as part of the World FinTech Report 2020 said their primary firm does not offer an integrated experience, making it impossible to access all of their accounts from a single platform.

Meanwhile, traditional banks are feeling the pressure of Neobanks and cloud-native fintech disruptors who have set a benchmark through their innovative digital banking solutions. As a result, customers now expect banks to offer the same seamless experience. Fintechs also have two major advantages over chartered banks – they aren't burdened with the costs of regulatory compliance, nor do they have to operate with restrictive IT legacy infrastructure.

From start to finish, the entire customer experience depends on an end-to-end value chain and interconnected systems. This includes everything from front-end customer engagement and convenient omnichannel interactions (in the branch, online, or via mobile phone) to middle and back-office operations, such as eKYC for risk management and compliance, loan or credit approval, underwriting, reporting, and analytics.

Modernizing piece-meal does not work. Instead, IT departments are forced to modernize systems by overhauling their entire legacy infrastructure. This is risky and can disrupt business continuity. Alternatively, businesses can adopt a phased modernization approach.

The journey can start with a low or non-disruptive modernization process. One way of achieving this is to use an in-memory data layer that's powered by Redis Enterprise to accelerate legacy systems. This will provide consumer-grade experiences by decoupling and modernizing monolithic systems to microservices-based architectures.

Below we'll show you how a modern layer provides banks with non-disruptive innovation and how it can overcome the issues with legacy systems. We'll then explore how a real-time data platform can maximize performance and help banks achieve digital transformation, without disruption, through a number of different use cases.

1 451 Research, a part of S&P Global Market Intelligence, *Financial organizations must prioritize CX and technology to drive loyalty in a digital-first landscape,* October 18, 2021

How Legacy Systems Impede Digital Transformation

Covid-19 has accelerated digital transformation roadmaps and continues to make this a high priority. However, the pandemic has also revealed the cracks in the way systems have been designed, often relating to the inadequacy of IT legacy systems and their lack of connectivity between existing IT systems, applications, and data.

IT legacy systems have prevented banks from evolving in tandem with the market's increasing shift towards digital technology. The ramifications from this are severe: delayed app rollouts, performance inefficiencies, and higher costs when delivering digital initiatives.

For traditional banks, legacy IT systems dependent on RDBMS on mainframes along with new database technology running on disk-based storage all cripple performance levels in so many ways. Legacy infrastructure is slow, cumbersome, and incompatible with today's fast-paced digi-sphere where success hinges on a firm's ability to guarantee real-time responsiveness.

Changes in customer expectations have shaped the digital experience and financial institutions face a myriad of challenges when trying to modernize their data layer.

Slow Performance

Legacy IT systems are slow and are considered prehistoric in the technological age. They were built in a time where banks operated only during branch hours and as a result, used batch processing to process high volumes of structured data overnight. Today banks operate 24/7 and are drowning in data.

Now legacy IT systems have to process, filter, and streamline an unprecedented amount of semistructured and unstructured data in real-time to provide users with an accessible and seamless banking experience. Relational databases were not built for this, nor were they built to provide large-scale, simultaneous access to real-time data.

Consequently, banks are struggling to leverage the treasure troves of insights they're sitting on to capitalize on business opportunities and provide digital banking products catered to the modern consumer. This is epitomized in a MuleSoft survey where 82% of LOB managers state that they need quick and easy access to data to do their job effectively and remain productive (ZDNet, 2020).

Fragmented Data Silos

Over the years, many traditional banks have built or acquired separate point solutions to meet the needs of the different businesses and customer interaction channels. This has led to the fragmentation of data silos, making data synchronization even more challenging.

Having disparate databases for each business area and channel prevents banks from gaining 360 views of customer data in real time, leading to an incoherent banking experience when users have multiple products or move between channels.

Customer insights are locked within these data silos and the process of extracting them is tedious, slow, and can hamper the overall quality of data. Banks therefore find it difficult to utilize these insights to create tailored digital solutions and personalized recommendations based on market and customer behavior.

The same survey suggests that banks are recognizing the negative impact that siloed data systems has on performance, with over 59% of LOB managers stating that data silos hamper business growth (ZDNet, 2020).

Inefficient

Legacy systems hinder cost efficiency due to expensive technologies and talent shortages. The MuleSoft survey indicated that 63% believe their IT organization lacks the skills and technology to keep up with the pace of digital projects (ZDNet, 2020). The mainframe systems used by banks are prehistoric in technological terms and there aren't many specialists available with the right skill set to maintain or update them.

In most cases, these mainframe databases cannot be replaced in the short term and require the addition of a modern operational data layer to improve efficiency and performance without disruption.

Not Resilient

Applications today must be highly-responsive and resilient. Any downtime can have catastrophic short and long-term consequences that begin with a hampered user experience and end with a drop in the number of customers.

Databases must guarantee high availability that consistently meets user expectations during normal and peak times. However, the growing digital demands are fast exceeding the capabilities of banks' legacy systems which are leading to database failures and application downtime.

Through periods of downtime, banks will experience hampered productivity, potential data loss, damaged brand value, a loss in revenue, and much more. However, eventual database failures are inevitable, placing greater importance on databases to be resilient and minimize application downtime.

Redis Enterprise – The Modern Data Platform

Redis Enterprise is an in memory NoSQL data platform that powers the next generation of financial services applications. It offers multiple data models with real-time performance and reliability across any environment to help organizations get the most value out of their data layer. There are 4 key capabilities that make Redis Enterprise the ideal modern data platform:

1. Sub-millisecond responses delivering instant customer and client experiences

Redis Enterprise delivers high throughput and sub-millisecond latency enabling the instant and fully interactive experiences customers and clients expect. The linear scalability of Redis Enterprise can offload data from legacy systems and speed-up workflows, This enables enterprises to maximize the customer experience across channels and launch innovative digital products.

2. Freedom and flexibility to develop modern applications with a unified data layer

With Redis Enterprise, organizations have the flexibility to build and scale modern applications without compromising performance or availability. Redis supports a wide range of use cases with purpose-built data models including real-time search, graph processing, event streaming, timeseries, probabilistic filters, and AI/ML model serving that can be applied to use cases like fraud detection, wealth management, quantitative trading, and other innovative solutions. With direct inter-model communication and single dataset copy with consistent access, Redis provides app developers a unified data layer, no need to have specialized databases for each microservice or application. Redis Enterprise is a simple to use modern data stack that enables enterprises to build apps faster and more efficiently.

3. Efficiency and faster time-to-market with managed service in the cloud

Redis Enterprise offers unmatched flexibility with deployment choices across on-premises, multi-cloud, and hybrid cloud environments. Since some legacy systems and data cannot be moved to the cloud due to technical and regulatory reasons, the ability of Redis to work seamlessly across data centers and public cloud enables a more phased approach to app and data modernization. With fully managed service offerings on Google Cloud, Microsoft Azure, and AWS, financial institutions can migrate to a more flexible and scalable data architecture that improves development efficiency, enables faster time-to-market, strengthens compliance, and lowers TCO.

4. Enterprise-hardened and resiliency for mission critical applications

Redis Enterprise employs a shared-nothing cluster architecture that ensures fault tolerance at all levels with automated failover, as well as tunable persistence and disaster recovery. Active-Active Geo-Distribution enables global distribution with local latency for applications. A-A ensures resiliency and business continuity with a 99.999% uptime SLA. Enterprise-level multi-layer security is achieved by ensuring production data is isolated from administrative access and providing role-based access-control, authentication, authorization, and encryption.

Transformative Use Cases

Redis has partnered with leading financial services organizations to implement modern data technologies to accelerate the modernization and digital transformation process without significant capital and time investment.

The below typical use cases fall into three broad business domain areas: customer experience, risk management (including fraud detection), and analytics/reporting. These domain areas have the common need for real-time data and support for multi-structured data models.

Customer/Client Experience

Traditional banks today face increasing competition from Neobanks determined to grow their market share. Tech giants such as Apple, Amazon, and Google are breaking into the finserv industry with their digital, mobile app, and cloud experience, tipping the scales in their favor as banks continue to struggle to maximize the user experience.

To level the playing field, traditional banks must modernize and improve their interactive customer apps by providing a personalized customer 360 omnichannel experience – a challenge that requires a complete revamp of their customer engagement strategy.

But should banks be able to provide a full-service mobile banking app, this has the potential to boost customer loyalty, encourage more spontaneous interactions and provide banks with a deeper insight into customer activity.

Developing a deeper understanding of consumer behavior is particularly significant because it empowers banks to market more dynamically to individual customers based on their interests. By investing in modernization, banks are able to deliver the level of customer experience and personalization that customers desire, thereby ensuring their ability to remain competitive and avoid business risk.

Instant Customer Experience with Real-Time User Profile Store

Overview

Digital banking is now the norm and consumers are using mobile devices to access their accounts daily. A top ten bank in North America wanted to maximize the user experience to beat their traditional and fintech competitors.

Their goal was to provide users with instant access to all account information since many of their customers have multiple products with them, such as savings, credit cards, loans, and investments. In addition, this bank wanted to provide tools and insights based on their transaction history to help their customers manage their finances effectively.

Challenge

The different account and transaction data associated with a specific customer was stored in a separate, siloed back-end disk-based RDBMS system. This technology was originally designed for batch processing, not online real-time access.

As a consequence, having to retrieve the needed data from each back-end system after a user logs in, as well as aggregating and presenting data in the proper format in the mobile app, led to long waiting times for customers.

This prevented them from offering a seamless, interactive user experience as well as meeting a crucial expectation - providing a consolidated 360 view of all customers' account information.

Solution

Redis Enterprise with the JSON module provides a real-time user profile and document store with native indexing, querying, and full-text search capabilities. It acts as a query accelerator for the bank's front-end applications.

During the app or website login process, the user device ID is identified and associated with the user profile. In the background, the app makes asynchronous calls to the slower back-end systems, retrieving and assembling the data in the JSON record object under the session GUID key and temporarily storing it in JSON.

Limited account value information is presented without any PII. By the time the customer is fully authenticated, all the details of all their accounts are ready to be viewed. Redis instantly provides customers and clients with consolidated account and balance information with no lags, freezes, or delays.

In addition, by loading the most recent 6 months of credit card transactions, a view of spending behavior and trends can be shown to the customer. As an example, by calculating the spending over that period of specific categories, the bank was able to highlight that there was a 20% increase in their grocery or gas station purchases compared to the previous year.

Machine learning algorithms could then be applied to these sets of transactions to provide users with personalized marketing promotions. This is a powerful way of tailoring the banking experience to the user's needs and wants, helping to build stronger relationships with customers and retain their loyalty.

REFERENCE ARCHITECTURE

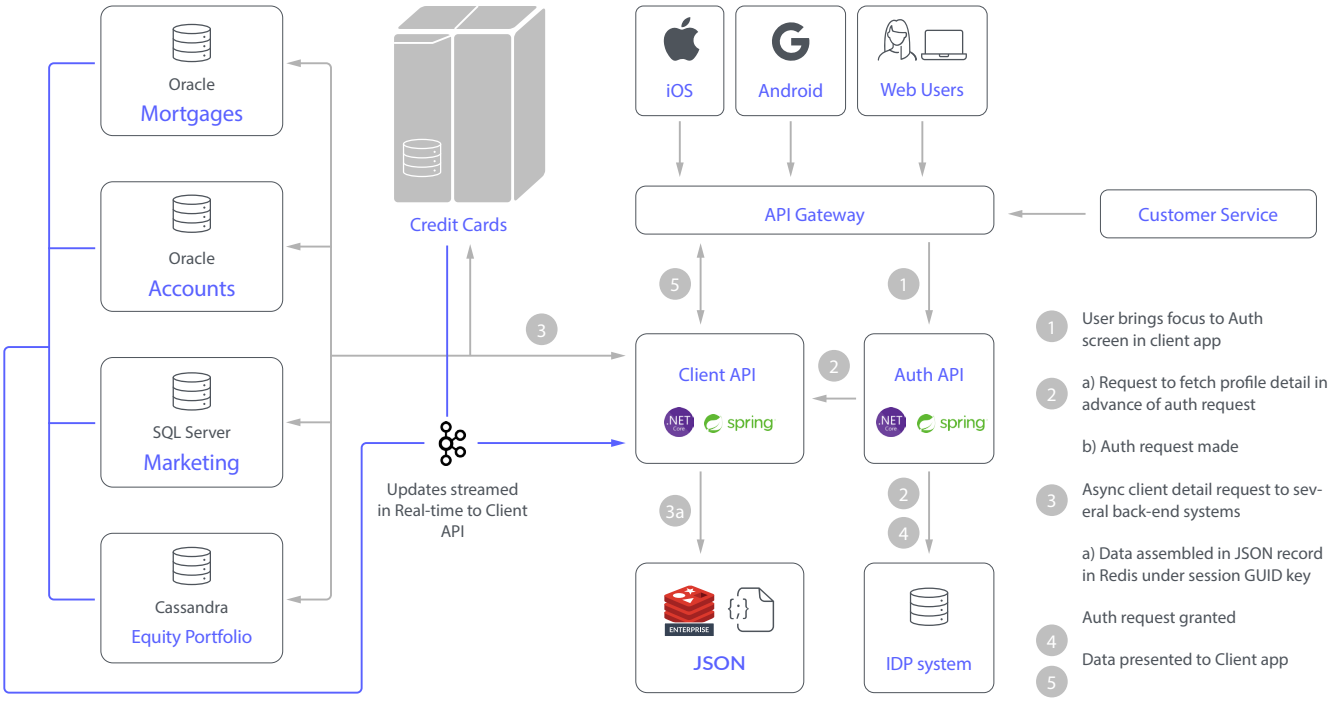

Fast Reference Data for Securities Trading

Overview

Securities trading is complex. A large security trading firm was rewriting its custom execution management system to facilitate and manage the execution of trade orders. Each order requires specific information to complete and settle the transaction. This could be, for example, reference data (e.g. security identifiers such as the CUSIP, ISIN, and SEDOL descriptors) as well as trade IDs. Reference data can be very dynamic for complex financial instruments which require them to always be available and updated frequently.

Challenges

Transactions, trade IDs, and reference data were stored in MongoDB and Kafka were used as part of the message broker and event streaming architecture in the new trading microservices-based app. Due to the slow performance of Mongo DB running on disk-based storage, Kafka was also used as a cache. However, this did not improve the performance since the additional roundtrip to Kafka reduced the number of trades by up to 95%.

Solution

Redis Enterprise's in-memory database was implemented for the specific reference data look-up and trade ID microservice due to its low latency, high throughput performance and availability. Redis replaced both Kafka and MongoDB as the caching solution, and reference database.

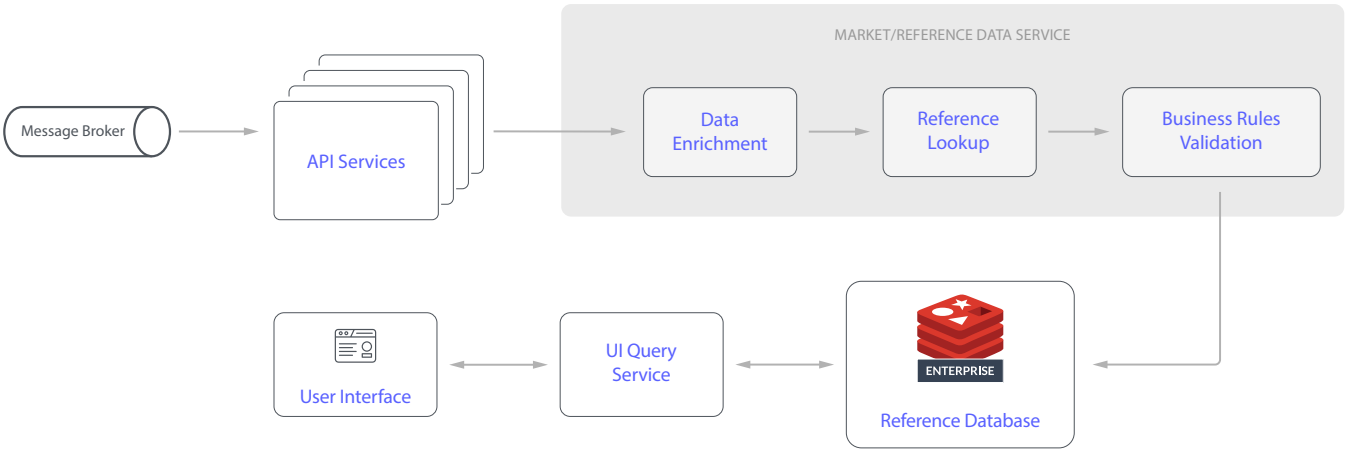

Risk Management

Many different types of risks pose a threat to financial services institutions, including operational, compliance, and financial risk. Operational risk refers to any risk incurred as a result of failure in people, internal processes, policies, and systems.

System failures involve critical service interruptions and security breaches. The 2008 financial crisis precipitated a wave of regulatory fines and enforcement actions on questionable sales practices, financial crimes, and employee misconduct.

At the same time, digitization, automation, remote work, and fintech partnerships have increased the reliance on software, IT systems, and the internet, creating new cyber risks and single points of failure.

Banks have focused on governance to mitigate these human-risk factors and internal process failures, putting in place elements such as loss-event reporting, risk control self-assessments (RCSAs), extensive control processes, and developing operational-risk capital models.

Rules-based fraud and trading alert detection systems have been leveraged by financial services organizations for quite some time now. Moreover, all of these tools have proven to be ineffective in proactively detecting cyber-risk, fraud, aspects of conduct risk, and other critical operational risk categories.

Studies from a McKinsey report on operational risk management highlight that these fraud-detection techniques can have false-positive rates of more than 90%.

As part of a Governance, Risk, and Compliance (GRC) framework, risk management is critical in today's world of cyber threats, identity theft, fraud, and automated financing – not to mention the inherent financial risk in asset management and investing.

Given the complexity and costs involved, a new approach is needed to detect these risks and issues. Banks and financial institutions can and should tap into the large repositories of structured, semistructured, and unstructured data to identify risk issues, move beyond reliance on self-assessments, subjective controls, and manual investigations.

Data access and analytics can transform risk management by helping financial services organizations move from qualitative, manual controls to data-driven, automated real-time detection. As an example, by mining sales and customer data, banks can detect potentially unauthorized sales or illegal practices.

At the heart of these predictive analytics tools is the use of machine learning models which can identify cyber threats, fraud, and even money laundering. Another example would be using natural-language processing to analyze calls, emails, surveys, and social media posts to identify spikes in risk topics raised by customers in real time, while complying with applicable privacy regulations.

Granular and Efficient Financial Risk Analysis

Overview

Effective financial risk management is fundamental to mitigating potential losses in all areas of capital market asset trading, which can involve significant investments in different asset classes, commodities, options, futures and more.

In this use case, one large financial services organization carried out daily risk assessments to reduce the level of risk for both themselves and for their clients. This is significant because having recurring risk calculations will indicate whether more liquidity is required to cover the risk.

Challenges

This financial firm wanted to improve the scalability, granularity, and frequency of their risk calculations as market volatility and their client investments grew. But the main problem was that the loading process of the config and market data to their risk engine calculation node was bottlenecked by both the NFS and local filesystems

The virtual machine hosting the risk engine node was also a bottleneck and therefore unreliable. This is problematic because each calculator compute node has to be large and of sufficient size to gather all of the market data and calculate risk results.

Consequently, this requirement meant that the entire workflow was unable to scale effectively for additional asset classes (groups of market data + risk calc nodes). A lack of resilience was also an issue, as each risk engine node cannot resume work for another node if one went down.

Solution

By using Redis Enterprise as a cache/temporary data store, the data was able to be loaded more efficiently into the risk engine calculator nodes. If at any point the firm identified a large asset class, they were able to dynamically allocate N number of nodes independently since they also had the engine configuration stored in Redis.

Thanks to Redis, they were able to create nodes for different asset classes. Should there be a node that's busy that day due to external events (e.g. geopolitical events) that affect that particular class, the financial firm would be able to spin up more risk engine calculator nodes.

The result is that these nodes are now able to be downsized since they no longer require a whole copy of data on the virtual machine heap. Smaller nodes that can scale out provide better resource utilization in their system and can run these risk calculations on a much more frequent basis each day then previously.

Scalability also improved significantly since the workflow can be easily duplicated for each asset class to improve overall risk calculations. Resilience was further enhanced as each risk engine node can resume work for another.

For the business, they were able to make risk positions more frequently available for themselves and their clients whilst reducing infrastructure costs. Now they're able to establish the liquidity risk earlier with great accuracy thereby freeing up additional capital to be reinvested elsewhere in the business.

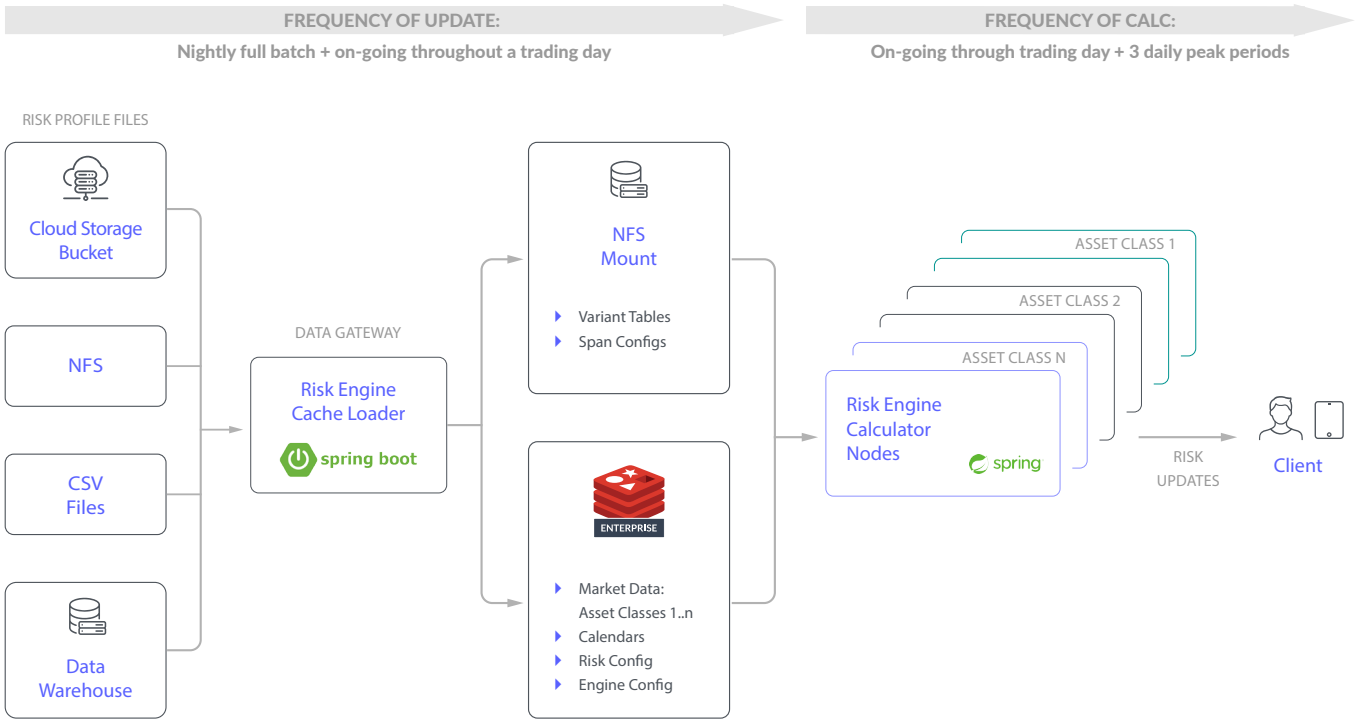

Reduced Costs With Efficient Case Management and Reporting

Overview

In an ideal world, fraud should be detected automatically and during the time when a transaction is taking place. However, even with the most sophisticated AI/ML algorithms, cases of fraud must still be reviewed and investigated manually.

There's a large US financial institution that employs investigators who use case management tools to carry out research and determine whether suspicious transactions are either fraudulent or related to money laundering. This company was using an in-house developed case management system that stores a large database of transactions and related data for investigators to sift through.

Challenge

Given the number of ongoing cases, the firm was looking for a solution that was able to improve the efficiency and productivity of their case investigators that had to do extensive searches of the data. Initially they tried using some traditional disk-based search solutions including ElasticSearch but these did not meet their performance requirements due to network latency as well as the need to encrypt files and data since it was in a persistent storage system. The encryption/decryption slowed down the search process considerably.

Solution

Redis Enterprise provided a fast in-memory search solution by using the Search and Query module to index the data and provide a hyper-efficient processing engine for searching discrete fields, terms across cases as well as discovering patterns that would indicate fraud.

Since both the data and attributes were in Redis' memory and non-persistent, it was able to remain de-tokenized and be ephemeral, just as it would be in the UI. This meant that no encryption was required and that searches could be performed in milliseconds.

To maintain the latest data, Redis Connect was used as the Change Data Capture (CDC) framework to extract the updates of the case data from the source relational database to Redis.

The Redis Enterprise solution along with the lightning-quick memory search capabilities of Search and Query boosted the productivity and efficiency levels of case investigators. For the financial institution, this led to a decrease in costs and time expenditure to close out cases.

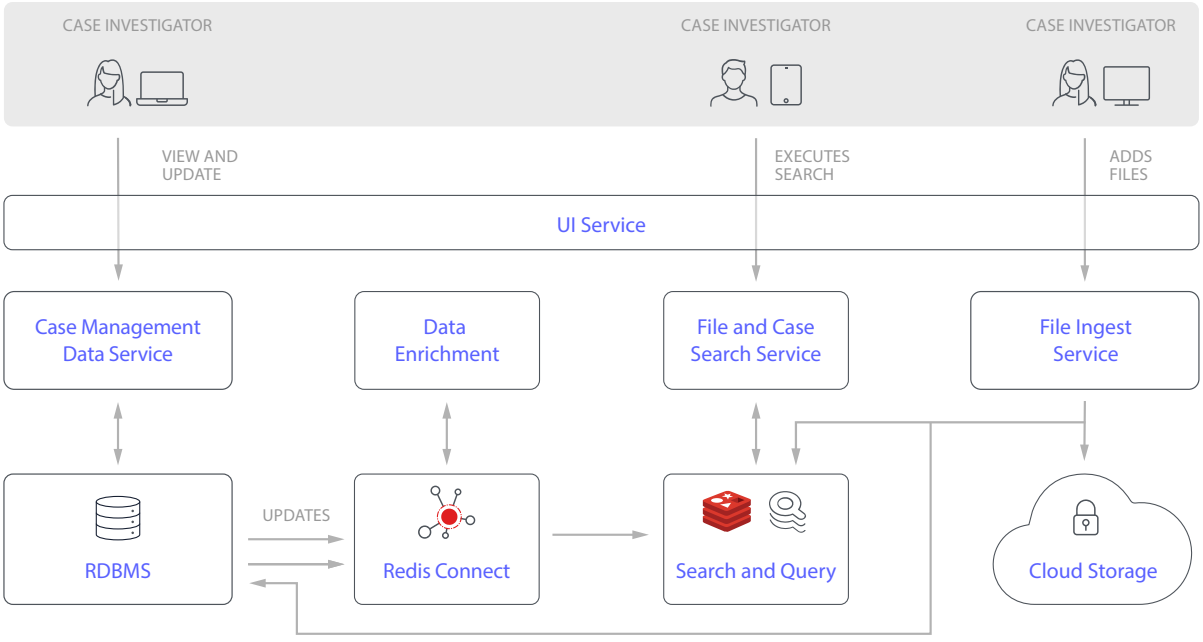

Enable Zero Trust with Granular Access Management

Overview

Cybersecurity at financial institutions and banks continues to be a priority and a large percentage of the IT budget is devoted to it, as security breaches and cyberattacks increase.

At the same time, the sharing of financial data outside the bank is also increasing due to the adoption of banking as a service, the expanded partner ecosystem, and compliance with Open Banking regulations.

In the midst of these changes, a leading North American bank wanted to differentiate themselves from their competitors by building a zero-trust OAuth-based identity management system that enables them to become the trusted identity provider.

What's important to note here is that in a zero trust model, every request is allowed or denied in real-time, whereas the typical OAuth2 implementation allows tokens to be available for up to fifteen minutes before they expire. An attack could occur within that fifteen-minute window.

Moreover, this bank also wanted to use an Attribute Based Access Control (ABAC) model for greater granular access to the APIs, data, and underlying infrastructure resources. As an example, for integration with a tax platform or fintech partner (Mint), the bank only needs to share specific data fields for specific sets of users and developers.

Challenge

With every access request needing to be verified and with an ABAC model, the granular Policy Decision Points (PDPs) required high performance and availability. This bank had to support global clients and developers who needed access at all times and locations. The complexity and high infrastructure cost to make traditional SQL relational databases high–performing and scalable as well as having the capacity to support global availability was not an option for the bank.

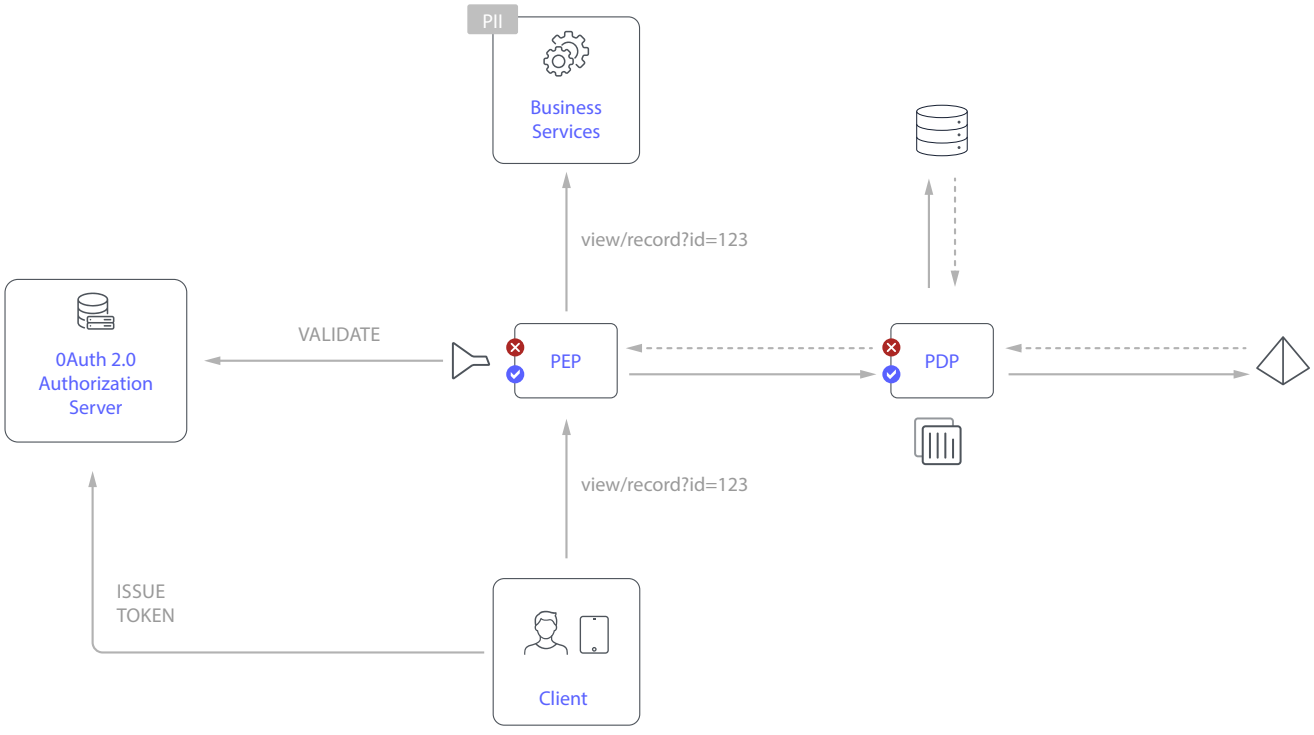

EXAMPLE OF OIDC/OAUTH BASED IDENTITY MANAGEMENT SYSTEM

Solution

By leveraging Redis Enterprise, this bank was able to overcome these challenges and provide global clients with real-time secured and granular access at any time and at any location. Below are some of the key features:

- Active-Active geo-distribution for high performance
- Globally active Authorization Server and ABAC primary database to store OAuth2 tokens, user profile/IDs, and general permissions

Redis can scale linearly as more APIs, clients/developers, and PDPs grow with an expanding number of apps and microservices. Each access hits a PDP for better security. Thanks to Redis, the bank became the trusted identity provider and enabled developers to innovate and leverage the data in both a secure and authorized way.

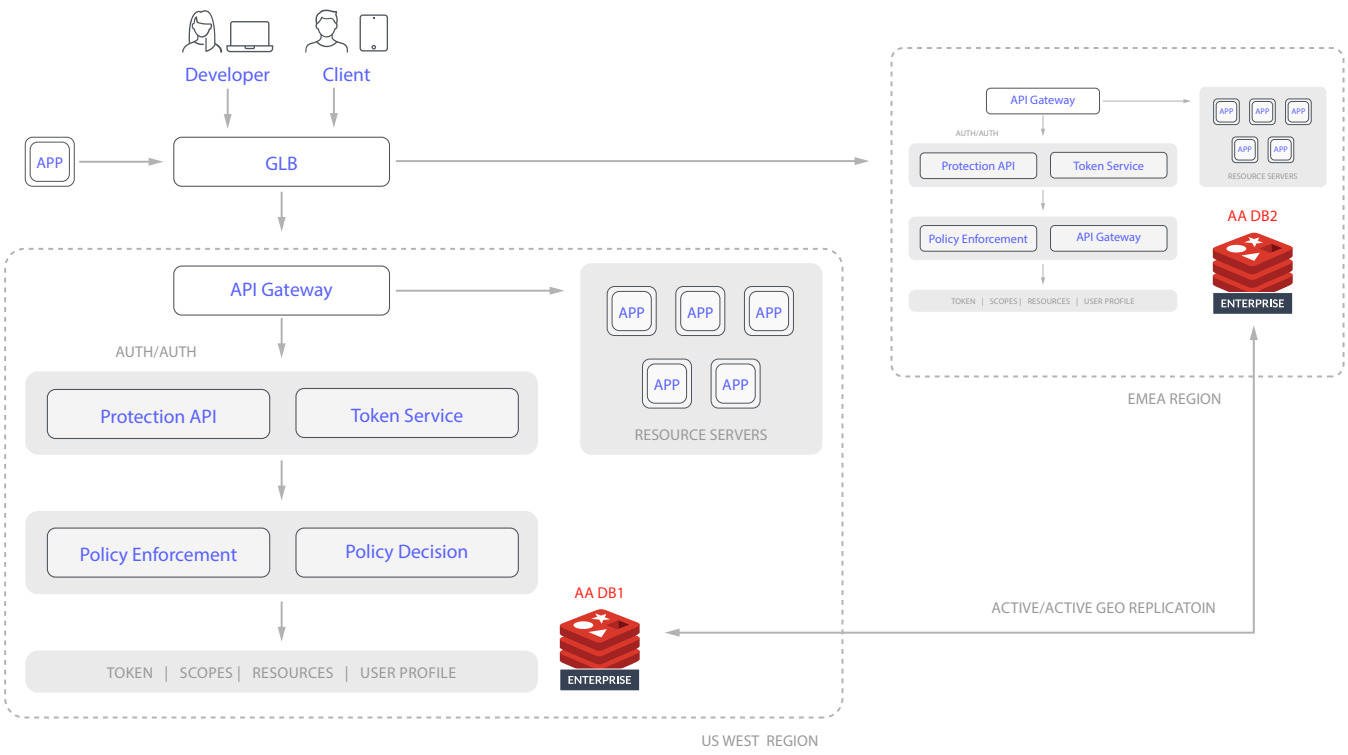

Fraud Detection Using AI/ML Models

Overview

The fraud epidemic is real. Last year fraud cost businesses worldwide $5.38 trillion, creating an erosion of trust and loyalty amongst consumers to banks. The increased complexity, volume, and speed of today's online transactions mean that organizations need to utilize more advanced fraud detection methods to suppress the ongoing onslaught and changing tactics of cybercriminals.

Rules-based engines have been used for decades to combat fraud. In today's fast-paced environment, they are rigid, resource-constrained, and time-consuming to adjust. These IT systems are effective in detecting simple, non-changing, known patterns, such as black lists or performing velocity checks.

However, rules-only systems aren't adept in distinguishing risk from normal behavior, leading to an increase in false positives which, in turn, damages the user experience. Take for example a British customer spending $400 at a hotel in Rio de Janeiro - is this a sign of a cloned card being used abroad, or is it simply an executive on a business trip?

Machine learning algorithms are far more proficient in making this differentiation by combining large volumes of different data sets. These can include user profiles, transaction patterns, behavioral attributes and more to build and train risk-based predictive models that can instantly detect fraud.

Moreover, AI and ML models do not follow a binary view of whether rules criteria are met and can dynamically make adjustments for normal behavior or predictable patterns such as seasonal fluctuations.

In today's financial services industry, where a variety of new digital multi-channel products are reaching the market faster than ever before, where systems are becoming vastly more interconnected, and where cybercriminal activity continues to increase YoY, AI and ML models have become vital to combating fraud.

However, while AI platforms such as Tensorflow and PyTorch are great for training models, they were not built to serve them collocated with stateful vector data. Their need to query vector data over the network significantly hampers their ability to provide real-time fraud detection in-line with transactions.

The solution to increasing the accuracy and detection speed is to co-locate the AI models with the contextual input/output tensor data in an in-memory database, such as Redis, which provides the low latency needed to meet the transaction SLA.

Redis Enterprise is also used as the AI/ML online feature store (see the next page for the definition of a feature store) due to its scalable distributed architecture that has minimal latency, high throughput, and high availability.

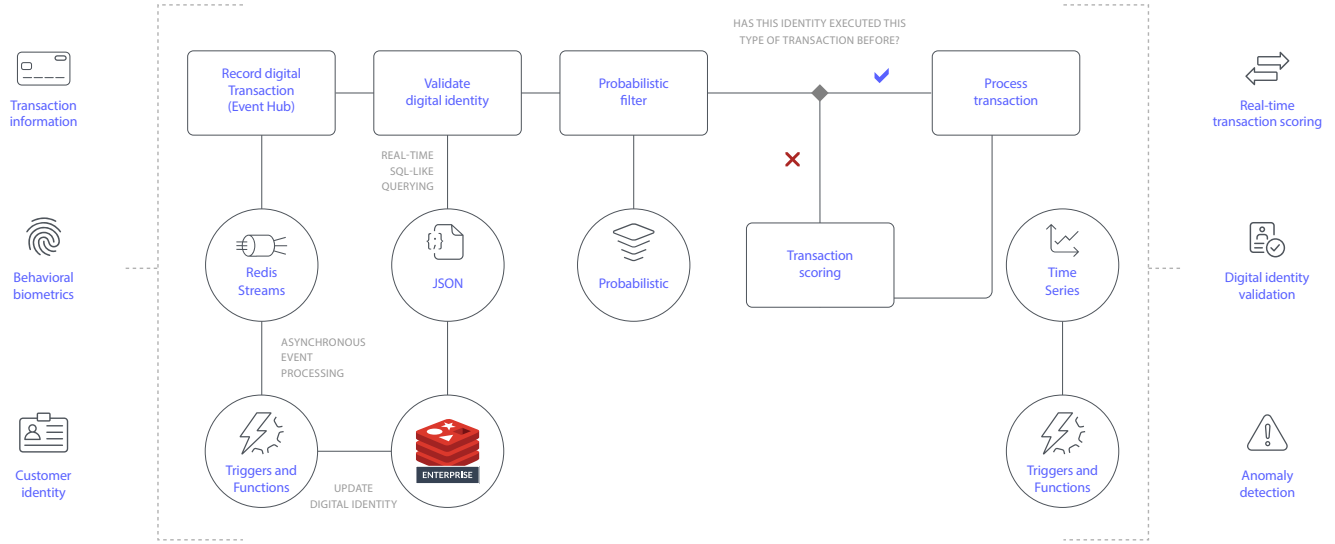

DEFINITION OF AI/ML FEATURE STORE

- Artificial intelligence (AI) is the practice of creating intelligent machines to simulate human thinking, e.g. general intelligence.
- Machine Learning (ML) is a subset of AI that enables machines to learn from data and make a prediction based on that data.
- Predictive applications rely on machine learning models trained on historical data to make predictions today about future events.
  - A Feature Store provides a repository of commonly-used "features" (data points) for data scientists, engineers and applications to work with. Features are a critical element of effective algorithms.
  - OFFLINE Feature Store: optimized for large datasets with complex point-in-time queries and used for model training
  - ONLINE Feature Store: optimized for low-latency queries of "smaller" datasets. Retrieve feature data as-of today and used for model inference
  - Examples of features for credit card fraud detection:
    - Customer
    - Transaction amount and date
    - Avg spend per transaction
    - Number of transactions in the last 5/60 minutes
    - Merchant
    - Location

Example of generic AI/ML Feature Store reference architecture:

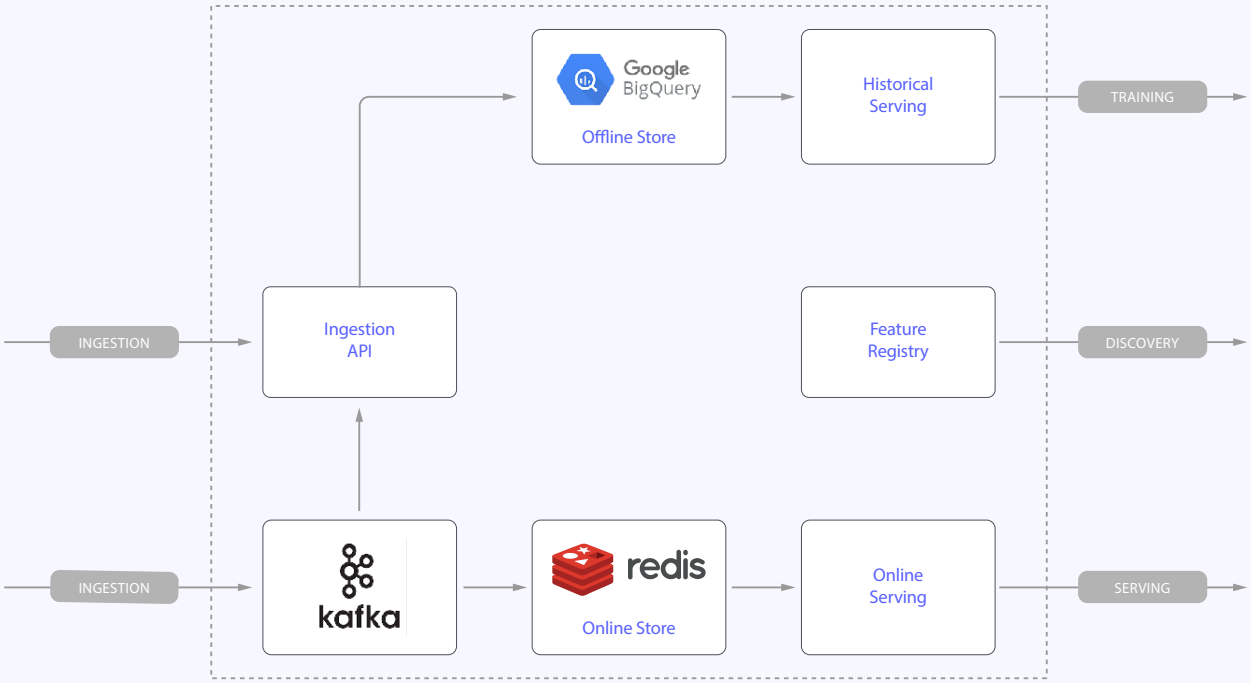

AI Online Feature Store

Overview

From speech recognition to email spam filtering through to self-driving cars, AI and ML technologies have been used to solve a myriad of different problems across different industries. The most pressing challenges facing banks and financial services organizations that require AI/ML algorithms include fraud detection, loan approvals, and risk scoring.

A large financial services company was strengthening their fraud detection system with a new custom real-time risk scoring microservice. The microservice determined a risk score which was fed to the fraud management solution that determined the final score and made the decision on whether the transaction was fraudulent or not.

Challenge

MongoDB was initially used as the online feature store to serve the online transaction features needed for the risk scoring model inferencing. However, MongoDB performance was insufficient for serving the features to enable real-time scoring on the incoming transactions.

Solution

Redis Enterprise replaced MongoDB as the online feature store and reduced the response time by 60X by leveraging its in-memory shared-nothing cluster architecture with a guaranteed 99.999% uptime SLA. Due to the unacceptable roundtrip latency to the Kafka cluster, this financial services company is also planning on replacing Kafka with Redis as the message broker solution.

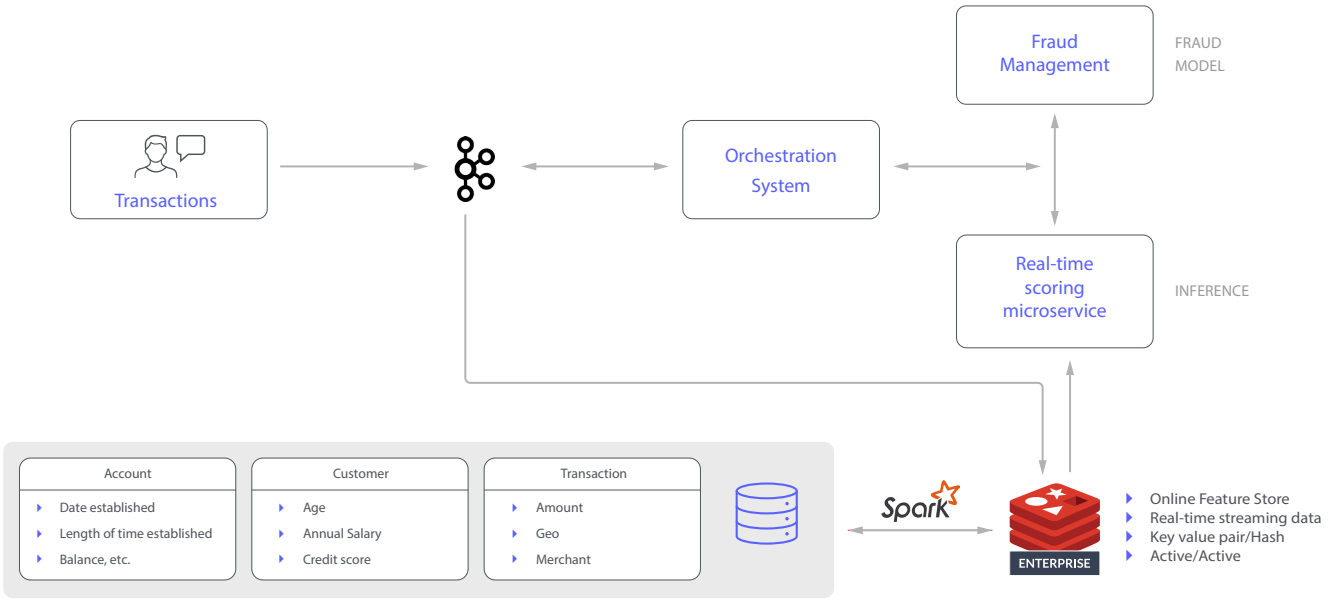

Analytics and Reporting

Financial services depend heavily on analytics for business insights. Many can act on data from months, weeks, or even days ago, but few can respond to updates or transaction streaming data that changes every minute or second.

As a consequence, many financial services organizations are stuck using batch processes, ETL tools, and data warehouse loads, and can't analyze semi-structured, unstructured, and geospatial data.

The type of data in the modern digital economy can change faster than legacy systems can handle. Batch processes work for some back office legacy workloads, but in many cases today, financial institutions need to analyze rapidly changing, multi-structured data in real time. To keep their finger on the pulse, it's mandatory for financial institutions to have access to real-time analytics.

While data warehouses have traditionally provided static reports, there is also an increasing move towards dynamic, customizable reporting. With Redis Enterprise as a cache in front of traditional disk based databases, and with support for most data structures providing the much-needed pre-sorting in-memory, Redis enables real-time analytics by providing dynamic querying over millions of records at sub-millisecond latencies.

By using Redis, you can analyze anything in real time. This includes chat history, ticks, social media posts, sales data, emails, satellite imagery, weather trends, and any other type of data that can be used to manage operational and financial risk.

In iterative processing scenarios, such as those with Apache Spark, Redis data structures accelerate processing and provide faster responses to analytical queries. The Spark-Redis connector package allows Spark to directly access Redis data structures for the most efficient in-memory data processing. Redis also offers a serving layer for Spark SQL and an accelerator for Spark processing.

A large multi-national bank uses Redis Enterprise Software to accelerate Big Data analytics in front of other disk based NoSQL databases, making analytic processing 45x faster.

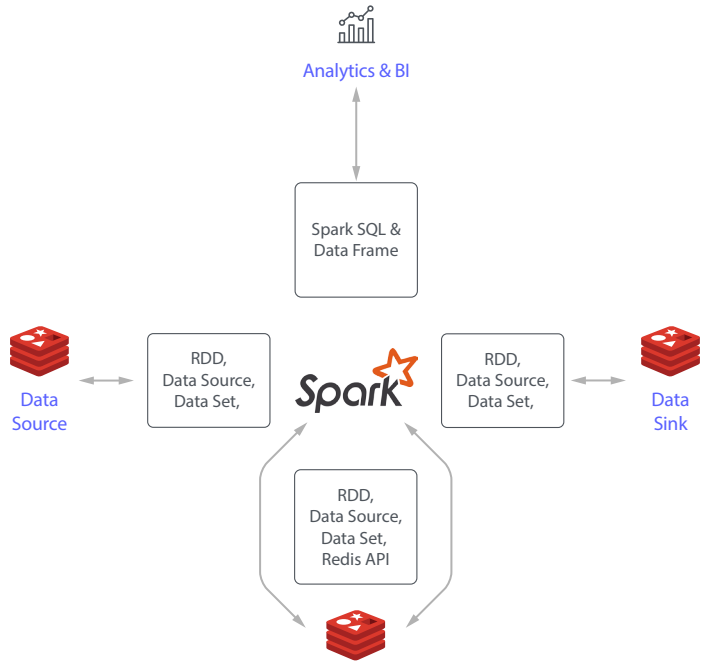

REDIS IS A DATA SOURCE, DATA SINK, ACCELERATOR AND SERVING LAYER

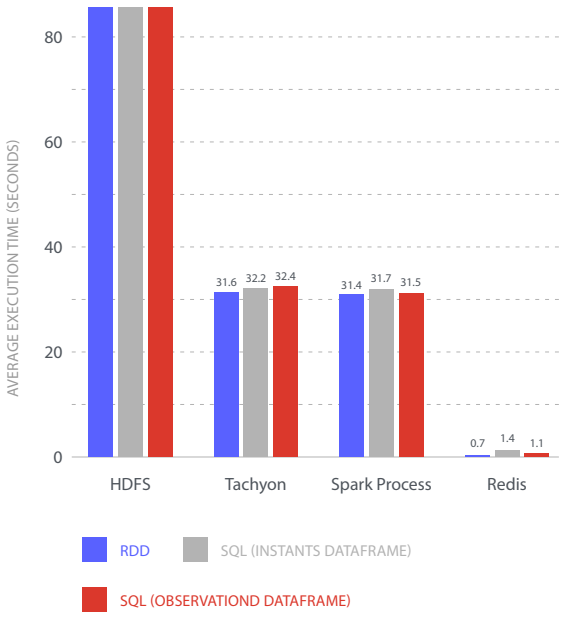

FASTER SPARK PROCESSING WITH REDIS BY UPTO 100 TIMES COMPARED TO HDFS AND 45 TIMES COMPARED TO TACHYON OR SPARK PROCESS MEMORY

Customer Success Stories

Here are a handful of stories of financial services companies successfully modernizing their data layer and enabling real-time data to power their businesses.

Deutsche Börse ensures low latency trade reporting services with Redis Enterprise

As an international exchange organization and innovative market infrastructure provider, Deutsche Börse Group offers its customers a wide range of products, services and technologies covering the entire value chain of financial markets. With around 6,000 employees, the company has its headquarters in the financial center of Frankfurt/Rhine-Main, as well as a strong global presence in Europe, Asia and the United States. In order to offer its customers a reporting solution, Deutsche Börse relies on rapid data reporting and processing.

Deutsche Börse turned to Redis Enterprise as an intelligent cache to rapidly process and organize data in order to fulfill this new requirement. As a result, Redis Enterprise has enabled Deutsche Börse to scale 20x the anticipated quote and trade data volume growth with no increase in latency.

Read the full customer story

Kenbi optimizes real-time payment approval with Redis Enterprise

Kenbi's innovative AI ecommerce payment platform provides a simple, straightforward solution that leverages the aligned interests of issuing banks and merchants to authorize more legitimate transactions. Kenbi's core solution relies heavily on data enrichment and risk management to help merchants and banks split the cost of risk to incentivize issuing banks to approve more transactions, Because time is of the essence in authorizing a sale before a shopper may choose to take their business elsewhere, Kenbi turned to Redis Enterprise Cloud on AWS. Beyond the performance Redis is known for, as a startup, Kenbi required a proven, fully managed service that offers guaranteed backup and redundancy to enable their team to focus on the challenge of optimizing the payment flow between e-commerce merchants and credit card-issuing banks.

Read the full customer story

Simility, a PayPal company, scales its real-time fraud detection cloud service

Simility is a PayPal service that combines machine learning and human analysis to provide a cloud based fraud detection service. With Redis Enterprise managing billions of transactions per day, Simility was able to deliver new application functionality 30% faster and improve overall performance by nearly 90%.

Read the full customer story

Summary

Many financial services organizations are feeling the pressure for digital transformation to address the need for instant customer experiences, real-time risk and fraud detection, and analytics. They must adopt a real-time modern data platform to meet the sky-high expectations of the modern consumer, while providing risk management leaders with the tools to combat operational, financial, and compliance risks.

The impact of IT legacy systems is apparent: slow performance speed, fragmented data silos, lack of application resilience, increased costs, and reduced innovation to respond to market changes as well as new regulations.

As a consequence, traditional banks are failing to provide a seamless omnichannel customer experience, are unable to detect fraud in real-time, and don't have the needed analytics to derive business insights. This will lead to frustrated users, damage brand reputation, and business losses. Fintechs, however, have no such problem and continue to innovate, adapt, and repeatedly set new benchmarks in digital banking.

Traditional banks must modernize or run the risk of being phased out by agile fintech disruptors who continue to exceed customer expectations. As demonstrated from the use cases, the financial institutions that implemented Redis Enterprise were able to seamlessly transition to an agile and flexible data architecture that supercharged performance levels in the key business areas.

With sub-millisecond latency, linear scalability, enterprise-hardened features, and multiple data models, Redis Enterprise is ideal for mission-critical financial applications. It ensures five-nines (99.999%) availability around the world with Active-Active Geo-Distribution across regions.

Redis Enterprise Cloud managed services are available in AWS, GCP, and Azure. Tiered storage options are available and provide you with a TCO by eliminating data center-related spending and improving IT productivity. This allows your organization to focus on rapid innovation rather than just keeping the lights on. In a recently commissioned TEI study conducted by Forrester Consulting on behalf of Redis, the analysis showed companies can achieve up to 350% ROI by using Redis Enterprise.

What do the experts say?

Forrester Total Economic Study (TEI) 2021

350%
ROI

$4.12m
Net Present Value (NPV)

$5.3m
Total Benefits Present Value

The Forrester Total Economic Study (TEI) is a comprehensive review of the value offered by certain businesses to help organizations understand the financial impact of investing in technology.

In this study, Forrester interviewed six customers who have used Redis Enterprise, combining the results from each experience to provide an aggregated financial analysis. Below is an overview of the results along with an excerpt from the report.

> *With the improved scalability of Redis Enterprise, interviewees avoided costs related to expanding their databases for both new and existing projects. Interviewees also improved their income by more quickly taking advantage of new business opportunities while recouping revenue previously lost to poor database performance.*

Want to download the full study?

Click here

Resources To Get Started

To get started, try building and testing a proof-of-concept real-time application. Deploy an in-memory multi-model database for new use cases and workloads. Get your feet wet now to make sure you're familiar with the key Redis Enterprise data structures and technologies that address some of the most pressing challenges in financial services today.

Key Resources:

- **Redis Developers Hub**
  The Home of Redis Developers
- **Video: Redis Stack Crash Course: How to Build Apps with Superpowers**
  Redis Stack turns Redis into a multi-model database, enabling you to build rich real-time applications.
- **Video: Declarative Caching with Redis**
  While caching is widely understood as a concept, not all queries should be cached. Declarative caching with Redis can provide finer control over which results are cached and for how long.
- **The New Stack: How Redis Simplifies Microservices Design Patterns**
  Learn more about how to scale and simplify microservice architectures with Redis.
- **The New Stack: How to Support Large-Scale Analytics with Probabilistic Data Structures in Redis**
  Probabilistic data structures can improve the efficiency of large scale systems that need to quickly filter through data, such as for fraud detection.

To learn more about how financial services firms are leveraging Redis Enterprise to build real-time FinServe applications, visit Redis Enterprise for Financial Services. To get started, try Redis Enterprise in the Cloud or download the Redis Enterprise Software for a free trial now.

About Redis

Data is the lifeline of every business, and Redis helps organizations reimagine how fast they can process, analyze, make predictions, and take action on the data they generate. Redis provides a competitive edge to any business by delivering open source and enterprise-grade data platforms to power applications that drive real-time experiences at any scale. Developers rely on Redis to build performance, scalability, reliability, and security into their applications.

Born in the cloud-native era, Redis uniquely enables users to unify data across multi-cloud, hybrid, and global applications to maximize business potential. Learn how Redis can give you this edge at redis.com.