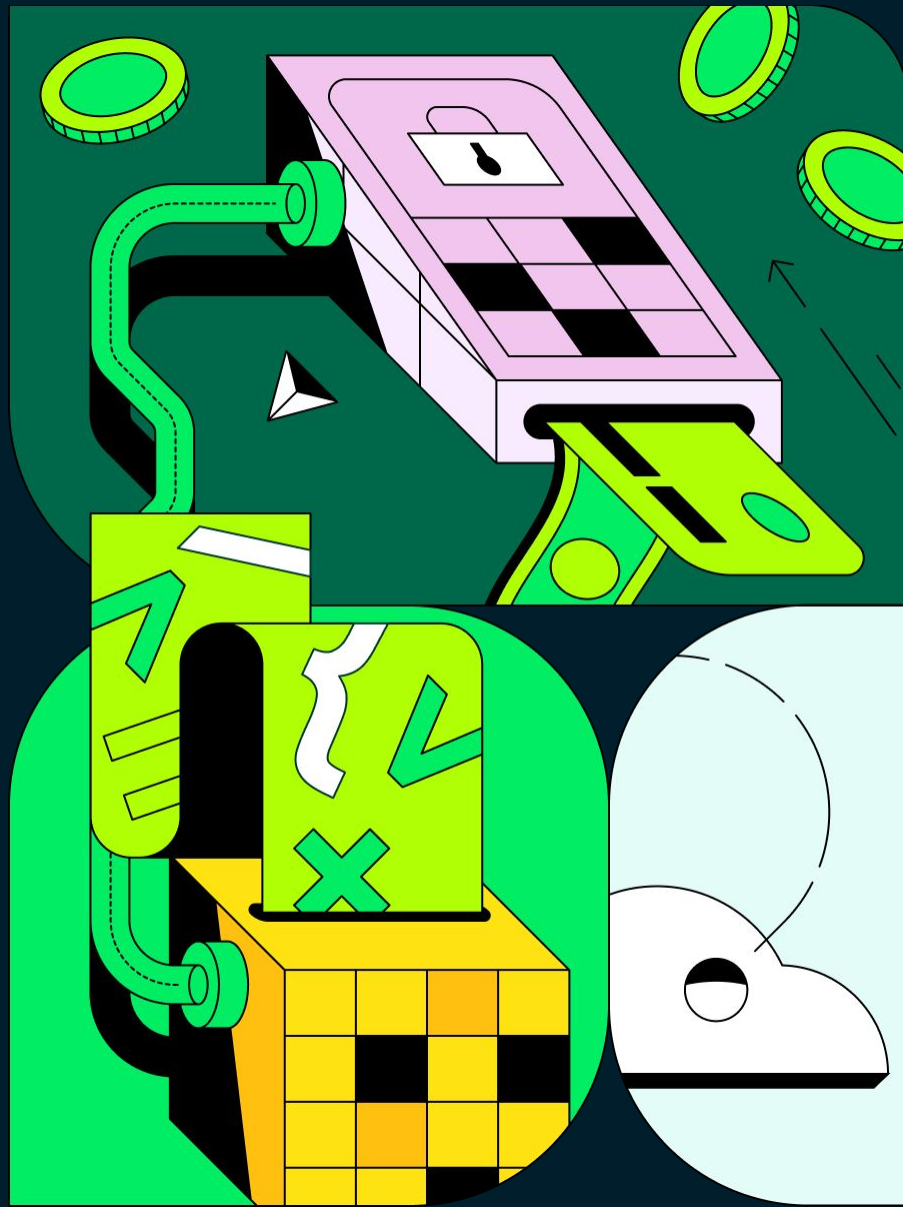




Atlas for Industries

Powering Innovation In Financial Services With Artificial Intelligence

Analyst firms, Celent and McKinsey, both identify financial services as one of the industries most likely to be transformed AI. While the initial benefits are likely to be process-related as firms slowly ramp up adoption, it won't be long before customer interactions are also completely re-invented by AI-driven experiences.



AI investments in the financial services industry



Already an integral part of the industry in areas such as risk, fraud, and compliance for example, the use of AI has been commonplace for years and continues to deepen, while there is a long (and growing) list of other use cases for which banks are also making use of AI.

AI technologies are used today to address a wide range of different workflows and customer-facing services from process automation and optimization in the middle and back office, to areas such as real-time risk and service personalization in the front office. However, it is to be expected that the benefits will be first in internal process optimizations and will only slowly start to be seen in direct customer engagement and interaction.

The conversation around generative AI became more nuanced through the latter part of 2023. This is understandable given the complexities of applying large language models (LLMs) to sensitive customer data, as well as broader regulatory concerns over the explainability of LLM outputs. That said, there are many areas in which gen AI is already being used to support advisors and relationship managers, with further innovation expected. According to Celent's [Harnessing the Benefits of AI in Payments](#) report, 58% of banks are evaluating or testing generative AI in some capacity (as shown below) while a further 23% have projects using this in their roadmap.

Interest in leveraging AI for banking use cases is expected to increase in the next few years.

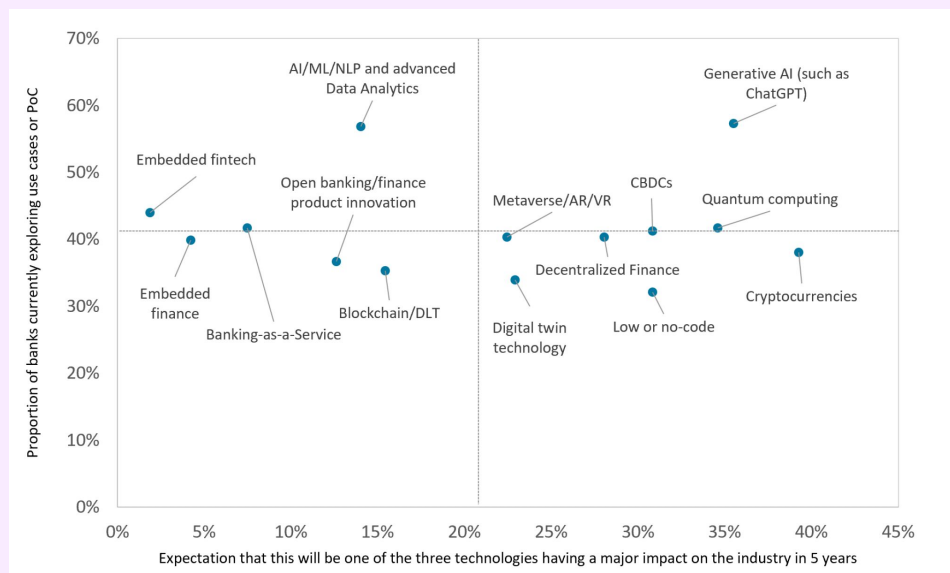


Figure 23: Bank tech exploration vs. expected impact

Addressing the challenges of AI in finance



While the industry has always had to deal with persistent issues like risk management and governance, adopting generative AI and machine learning introduces new challenges that AI specialists have always dealt with, like inherent biases and ethical concerns. One challenge that stands out for generative AI is hallucination—the generation of content that is not accurate, factual, or reflective of the real world.

Generative AI models, especially in natural language processing, might generate text that is coherent and contextually appropriate but lacks factual accuracy. This poses challenges, for example:

- **Misleading financial planning advice:** In financial advisory services, hallucinated information may result in misleading advice, leading to unexpected risks or missed opportunities.
- **Incorrect risk assessments for lending:** Inaccurate risk profiles may lead to poor risk assessments for loan applicants, which can cause financial institutions to approve a loan with a higher risk of default than the firm would normally accept.
- **Sensitive information in generated text:** When generating text, models may inadvertently include sensitive information from the training data. Adversaries can craft input prompts to coax the model into generating outputs that expose confidential details present in the training corpus.

A strategic and comprehensive approach encompassing aspects of technology, data, ethics, and organizational readiness is critical to overcoming these challenges:

- **Hallucination mitigation:** One promising strategy is using [retrieval augmented generation](#) (RAG) to mitigate hallucination in gen AI models—incorporating information retrieval mechanisms to enhance the generation process to ensure content is grounded in real-world knowledge. [Atlas Vector Search](#) is a popular mechanism to support RAG, which uses vector embeddings to retrieve relevant documents based on the input query.
- **Data quality and availability:** Take a step back before adopting AI to ensure the quality, relevance, and accuracy of data being used for AI training and decision-making can be accessed in real time.
- **AI education:** The key is to invest in training programs to address skill gaps, create a culture of learning and development, and promote awareness about vulnerabilities.
- **Develop new governance, frameworks, and controls:** Before going live, create safe and secure environments for testing.
- **Implement monitoring systems:** Measure and understand financial impacts, change impacts, scale, and complexity associated with AI.
- **Security and privacy:** Implement secure data access privileges and authentication measures to safeguard AI models and the data they rely on.

Relationship Management Support with Chat Bots



One key service that relationship managers provide to their private banking customers is aggregating and condensing information.

Because banks typically operate on fragmented infrastructure, this can require a lot of detailed knowledge about this infrastructure and how to source information such as:

- *When are the next coupon dates for bonds in the portfolio?*
- *What has been the cost of transactions for a given portfolio?*
- *What would be a summary of our latest research?*
- *Please generate a summary of my conversation with the client.*

Traditionally, relationship managers (RMs) faced challenges accessing and analyzing client portfolio data. Tasks like finding coupon dates were time-consuming and required manual effort. To streamline this process, RMs often requested new reports, but finding the right information could be difficult.

Generative AI offers a solution by automating data retrieval and analysis. Through RAG, AI can provide answers and explanations, and even execute queries based on diverse data sources, including portfolio information, user manuals, and customer data. RMs can interact with the AI to refine results and ensure accuracy.

MongoDB's platform is well suited for this application. It can handle various data types, including structured and unstructured information, and it provides efficient data search and processing tools. This enables the development of AI-powered applications that deliver timely insights to RMs.

There are different approaches of building a RAG-based chatbot application.

A simple approach would be like the one explained in our article [Taking RAG to Production with the MongoDB Documentation AI Chatbot](#) with an architecture diagram as follows.

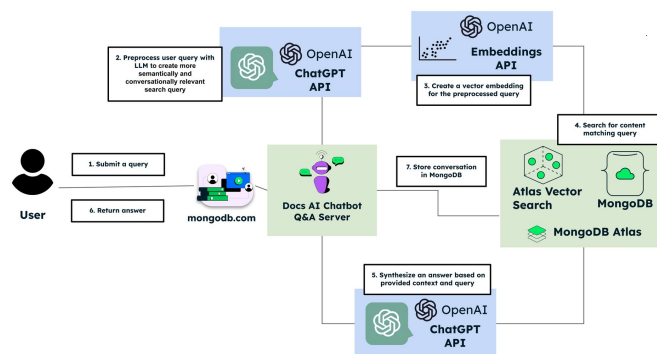


Figure 24: User query processing flow

The more advanced approach would be to leverage a composable RAG framework that is developed as part of the [MongoDB AI Applications Program \(MAAP\)](#).

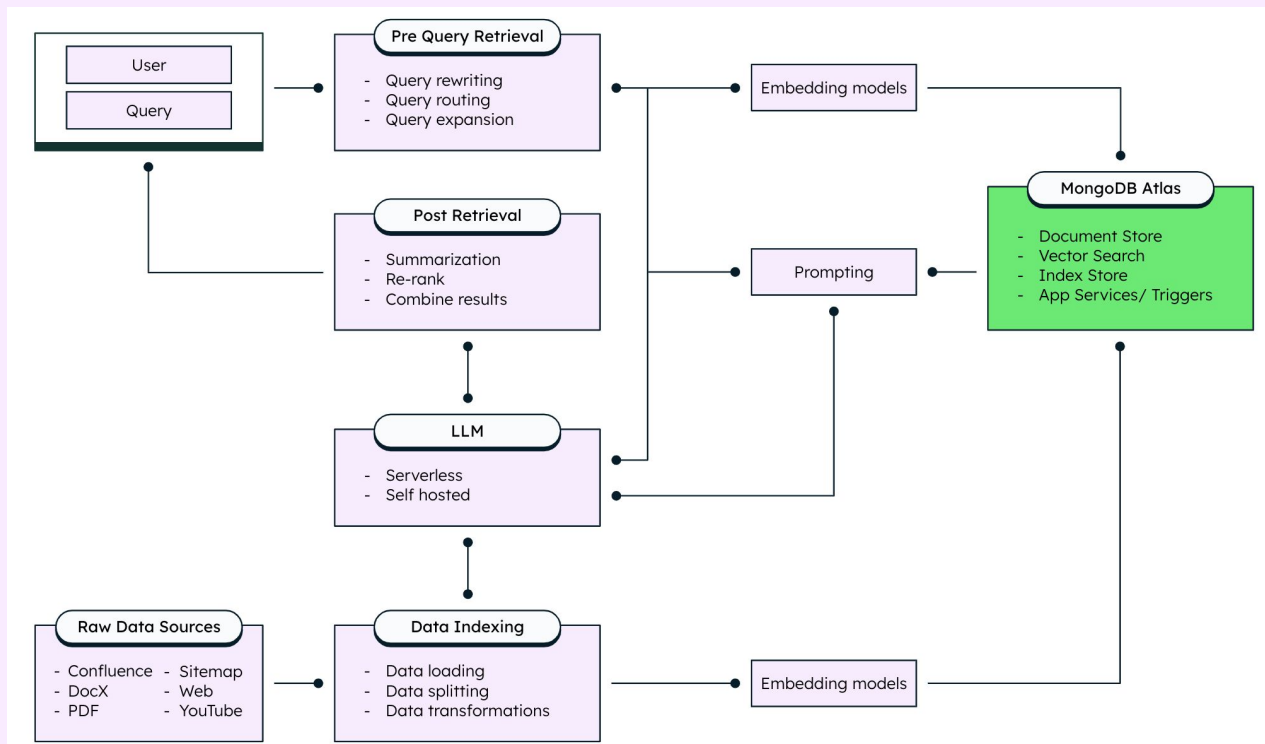


Figure 25: Query processing pipeline

Data loading from raw data sources

Begin by extracting data from databases or files. Examples of data can include customer profiles, transaction histories, and product information. The source of the data may be located in different databases or files and can be consolidated into MongoDB and serve as an operational data store to provide a customer 360-degree view.

Data indexing: chunking and vectorization

- **Data chunking** involves breaking down the documents (such as customer service records, emails, and transaction logs) into smaller, more manageable chunks (eg. individual interactions or sentences) and then creating an index that allows for quick searching and retrieval of relevant information.
- **Data vectorization** is implemented by using the embedding model (eg. Cohere, Anthropic, etc) to convert the above

mentioned text chunks into vector representations. The chunks can then be indexed using MongoDB vector indexing that will then enable an efficient retrieval of vectorized data using [Atlas Vector Search](#).

Pre-query Retrieval

Advanced RAG methods explore various transformations of user queries to enhance retrieval accuracy before it is passed to the LLM for response generation. As an example, If the RM asks for all recent fund transfers for a given customer, the gen AI application can pre-process the query by first checking if the RM has access privileges to the customer data. It can then refine the query to add a specific time period, and add the relevant documents retrieved from the vector search. This will augment the search, giving it context to optimize relevancy of the results.

Post retrieval

With the transformed query executed, the relevant documents will be retrieved (such as with regular, text, vector or even hybrid searches). The retrieved documents (such as the customer details and transactions) can be further filtered and reranked to provide an even more relevant set of documents to be used as a data augmentation when prompting the LLM later.

LLM

The LLM engine can be a single or a group of LLM engines deployed on an AI platform (eg. Fireworks.AI) to work together leveraging different specialization to provide a relevant, accurate and combined response to the users. Let us generalize and refer to these engines as **AI agents** so to not limit our imagination to just Language Models (be it large or small). These AI Agents not only can perform orchestration tasks but can agent can

also be trained to specialize in a specific domain for accuracy and efficiency. For example in the context of financial advisory support, an agent could specialized in Natural Language Understanding of financial market information, another AI Agent can specialist in Risk Assessment for evaluating client and product risk profiles while yet another can specialized in Compliance for regulatory adherence to ensure the investment decisions especially for corporate clients are also compliant. Together they can be “composed” using orchestration tools (eg. Langchain and LLamaIndex) to support the tasks required for the generative AI interactions.



Risk Management and Regulatory Compliance

Risk & Fraud Prevention

Banks are tasked not only with safeguarding customer assets but also with [detecting fraud](#), verifying customer identities (KYC), supporting sanctions regimes, and preventing various illegal activities (AML). The challenge is magnified by the sheer volume and complexity of regulations, making the integration of new rules into bank infrastructure costly, time-consuming, and often inadequate. ***For instance, when the EU's Fifth Anti-Money Laundering Directive was implemented, it broadened regulations to cover virtual currencies and prepaid cards.*** Banks had to swiftly update their onboarding processes and software, train staff, and possibly update their customer interfaces to comply with these new requirements.

AI offers a transformative approach to fraud detection and risk management by automating the interpretation of regulations, supporting data cleansing, and enhancing the efficacy of surveillance systems. Unlike static, rules-based frameworks that may miss or misidentify fraud due to narrow scope or limited data, AI can adaptively learn and analyze vast datasets to identify suspicious activities more accurately. Machine learning, in

particular, has shown promise in trade surveillance, offering a more dynamic and comprehensive approach to fraud prevention.

Modernizing fraud systems improves detection accuracy, increases flexibility and scalability, enhances security, and reduces the cost of operation. MongoDB's developer data platform is a great fit for modern fraud prevention systems due to:

Data visibility: MongoDB's document model incorporates any kind of data no matter how often it changes, allowing you to create a holistic picture of customers to better predict transaction anomalies in real time. As well as process large amounts of data and analyze in real-time to identify suspicious activity. [Change Streams in MongoDB](#) enable the monitoring of database changes in real time meaning immediate detection of suspicious activity or transactions. By receiving instant notifications, fraud prevention teams can quickly respond to potential threats.

Real-time data processing: By leveraging MongoDB as the [operational data store](#) (ODS) you can handle large amounts of data to effortlessly capture, store, and

process high-volume transactional data in real-time. [MongoDB's aggregation framework](#) allows for complex data processing and analytics directly within the database, crucial for implementing sophisticated, rules-based fraud detection algorithms.

Integrating with modern AI/ML

technologies: MongoDB's flexible data model makes it easier to integrate with AI/ML platforms to develop models for real-time, predictive fraud prevention that are trained on the most accurate and relevant data available.

Regulatory compliance: Compliance with regulations like PCI, GDPR, CCPA, and PSD2 is crucial in the finance sector. [MongoDB's security features](#) ensure data handling meets these regulatory standards. Enhanced authentication and role-based authorization controls further fortify the system.

Besides the use of AI/ML techniques, there is the option of leveraging vector data store capabilities such as [MongoDB Atlas Vector Search](#) to perform fraud and money laundering detection. This can also be used in conjunction with the AI/ML approach to create a hybrid implementation to further enhance the accuracy of detection and minimize false positives. [Learn More](#)

Below is an [example architecture of an ML-based fraud solution](#) leveraging MongoDB and Databricks. Key features include data completeness through integration with external sources, real-time processing, AI/ML modeling to identify potential fraud patterns, real-time monitoring for instant analysis, and model observability for full data visibility into fraud behaviors. [Learn more](#)

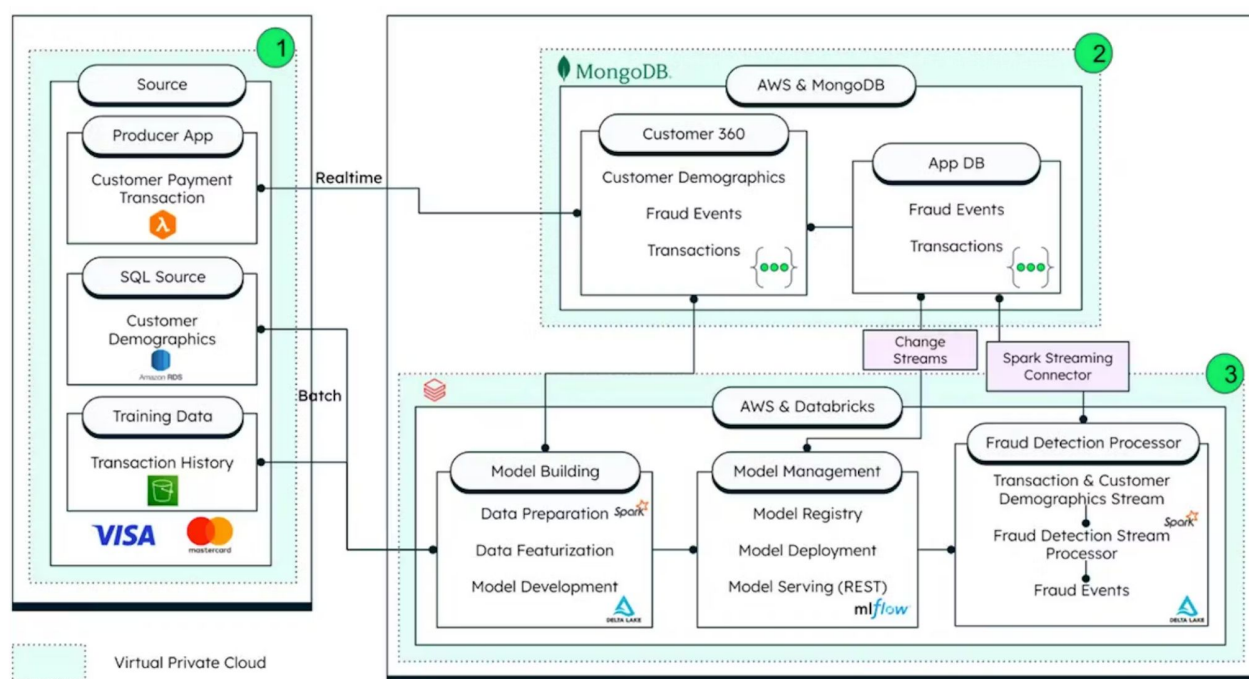


Figure 26: Example fraud detection solution leveraging MongoDB

Financial Document Search and Summarization

Financial institutions, encompassing both retail banks and capital market firms, handle a broad spectrum of documents critical to their operations. Retail banks focus on contracts, policies, credit memos, underwriting documents, and regulatory filings, which are pivotal for daily banking services. On the other hand, capital market firms delve into company filings, transcripts, reports, and intricate data sets to grasp global market dynamics and risk assessments.

These documents often arrive in unstructured formats, presenting challenges in efficiently locating and synthesizing the necessary information. While retail banks aim to streamline customer and internal operations, capital market firms prioritize the rapid and effective analysis of diverse data to inform their investment strategies. Both retail banks and capital market firms allocate considerable time to searching for and condensing information from documents internally, resulting in reduced direct engagement with their clients.

Generative AI can streamline the process of finding and integrating information from documents by using NLP and machine learning to understand and summarize content. This reduces the need for manual searches, allowing bank staff to access relevant information more quickly.

MongoDB can store vast amounts of both live and historical data, regardless of its format, which is typically needed for AI applications. It offers vector search capabilities essential for retrieval-augmented generation (RAG). MongoDB supports transactions, ensuring data accuracy and consistency for AI model retraining with live data. It facilitates data access for both deterministic algorithms and AI-driven rules through a single interface. MongoDB boasts a strong [partnership ecosystem](#), including companies like Radiant AI and Mistral LLM, to speed up solution development.

Discover more about MongoDB's full full-text search and vector search capabilities.

To achieve search and summarization using RAG, a conceptual architecture leveraging MongoDB and generative AI can be used.

This architecture aims to centralize document storage, enhance search capabilities, and automate the extraction and synthesis of information. The following diagram illustrates this conceptual architecture.

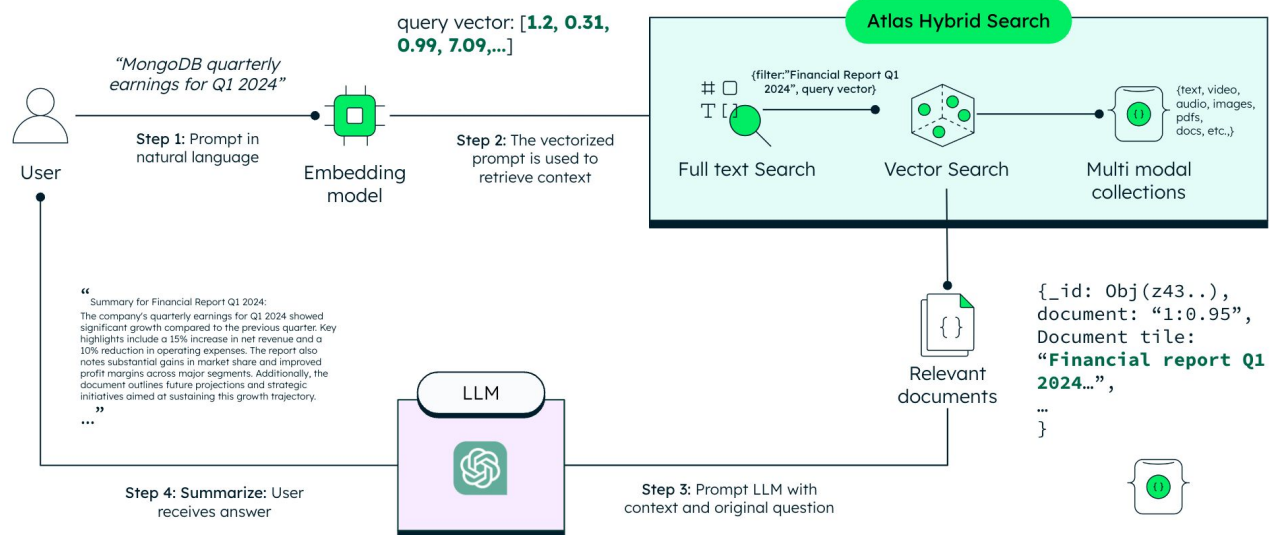


Figure 27: Atlas hybrid search workflow: query embedding, document retrieval, summarization, user presentation

ESG Analysis

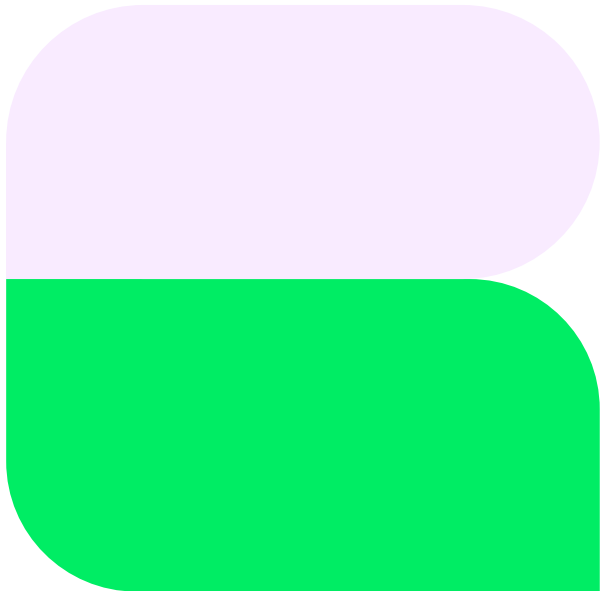
The profound impact of environmental, social, and governance (ESG) principles is evident, driven by regulatory changes, especially in Europe, compelling financial institutions to integrate ESG into investment and lending decisions. Regulations such as the [EU Sustainable Finance Disclosure Regulation \(SFDR\)](#) and the [EU Taxonomy Regulation](#) are examples of such directives that require financial institutions to consider environmental sustainability in their operations and investment products. Investors' demand for sustainable options has surged, leading to increased ESG-focused funds. The regulatory and commercial requirements in turn, drive

banks to also improve their [green lending practices](#). This shift is strategic for financial institutions, attracting clients, managing risks, and creating long-term value.


However, financial institutions face many challenges in managing different aspects of improving their ESG analysis. The key challenges include defining and aligning standards and processes, and managing the flood of rapidly changing and varied data to be included for ESG analysis purposes.

AI can help to address these key challenges in not only an automatic but also adaptive manner via techniques like machine learning. Financial institutions and ESG solution providers have already leveraged AI to extract insights from corporate reports, social media, and environmental data, improving the accuracy and depth of ESG analysis. As the market demands a more sustainable and equitable society, predictive AI combined with generative AI can also help to [reduce bias in lending](#) to create a fairer and more inclusive financing while improving the predictive powers. The power of AI can help facilitate the development of sophisticated sustainability models and strategies, marking a leap forward in integrating ESG into broader financial and corporate practices.

MongoDB's dynamic architecture revolutionizes [ESG data management](#), handling semi-structured and unstructured data. Its flexible schema nature allows the adaptation of data models as ESG strategies evolve. Advanced text search capabilities efficiently analyze vast semi-structured data for informed ESG reporting. Support for vector search enriches ESG analysis with multimedia content insights.



Incorporating LLMs enhances MongoDB's capacity to process ESG textual content, automating sentiment extraction, summarization, and trend identification. Combining LLMs with vector data management capabilities, generative AI applications can be created to interpret the complex and evolving sustainability taxonomy and guide the investment and financing processes in a compliant manner. This AI-driven approach, supported by MongoDB's robust data management, offers a sophisticated means of analyzing extensive narrative data in ESG reporting.



Furthermore, MongoDB supports geospatial and network graph analytics, providing a powerful combination of analytics to identify the physical risks associated with climate change (e.g., floods, wildfires) to assets financed by banks or investment firms and for assessing supply chain impacts of the climate risks. The risk analytics can then enable targeted strategies for risk mitigation and supply chain resilience.

MongoDB's value extends beyond ESG data management, accelerating productivity for developers and data science teams. Its intuitive data model, analytical tools, and AI integrations streamline the development and deployment of data-driven applications, making MongoDB pivotal for organizations advancing their ESG agendas efficiently.

Below is a diagram of an enterprise ESG solution architecture with the boxes labeled with a leaf where MongoDB can be deployed to support the ESG data analytics services.

Learn how MongoDB makes it incredibly easy to collect, analyze, and visualize ESG data in real time.

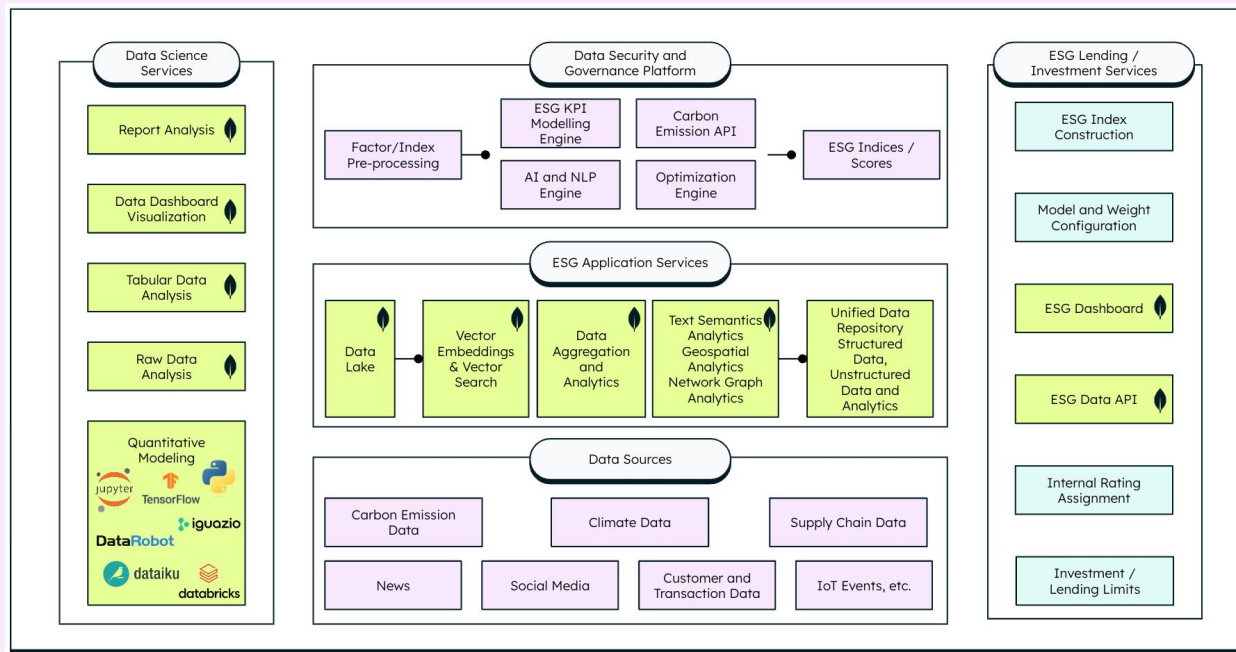


Figure 28: Blueprint for Enterprise ESG Solution Architecture Using MongoDB

Transforming Credit Scoring with AI

The convergence of alternative data, artificial intelligence, and generative AI is reshaping the foundations of credit scoring, marking a pivotal moment in the financial industry. The challenges of traditional models are being overcome by adopting alternative credit scoring methods, offering a more inclusive and nuanced assessment. Generative AI, while introducing the potential challenge of hallucination, represents the forefront of innovation, not only revolutionizing

technological capabilities but fundamentally redefining how credit is evaluated, fostering a new era of financial inclusivity, efficiency, and fairness.

The use of artificial intelligence, in particular generative artificial intelligence, as an alternative method to credit scoring has emerged as a transformative force to address the challenges of traditional credit scoring methods for several reasons.

Alternative data analysis: Unlike traditional models that rely on predefined rules and historical credit data, AI models can process a myriad of information, including alternative data, such as utility payments and rental history, to create a more comprehensive assessment of an individual's creditworthiness, ensuring that a broader range of financial behaviors is considered.

AI offers unparalleled adaptability: As economic conditions change and consumer behaviors evolve, AI-powered models can quickly adjust and learn from new data. This continuous learning ensures that credit scoring remains relevant and effective in ever-changing financial landscapes.

Fraud detection: AI algorithms can detect fraudulent behavior by identifying anomalies and suspicious patterns in credit applications and transaction data.

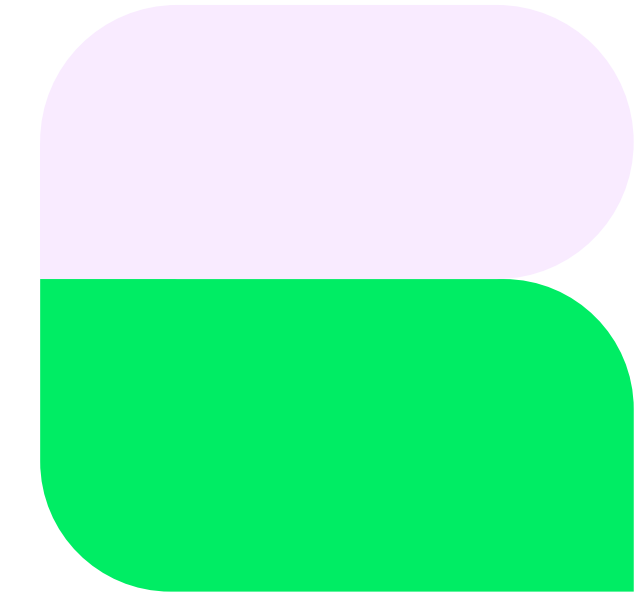
Predictive analysis: AI algorithms, particularly ML techniques, can be used to build predictive models that identify patterns and correlations in historical credit data, forecasting the greater likelihood of loan defaults.

Behavioral analysis: Many individuals encounter hurdles in the form of limited or nonexistent credit history, making it difficult to prove their creditworthiness due to the lack of historical data. AI algorithms can analyze behavioral data sets to understand financial habits and risk propensity. By monitoring real-time financial behavior, AI models can provide dynamic credit scores that reflect current risk profiles.

Lack of clarity in rejection reasons: A lack of transparency in rejection reasons leaves applicants in the dark, making it difficult for them to address the root cause and enhance their creditworthiness for future applications.

By harnessing the power of artificial intelligence, lenders can make more informed lending decisions, expand access to credit, and better serve consumers (especially those with limited credit history). However, to mitigate potential biases and ensure consumer trust, it's crucial to ensure transparency, fairness, and regulatory compliance when deploying artificial intelligence in credit scoring.

The most common objections from banks to not using AI in credit scoring are transparency and explainability in credit decisions. The inherent complexity of some AI models, especially deep learning algorithms, may lead to challenges in providing clear explanations for credit decisions. Fortunately, the transparency and interpretability of AI models have seen significant advancements. Techniques like [SHapley Additive exPlanations \(SHAP\) values](#) and [Local Interpretable Model-Agnostic Explanations \(LIME\) plots](#) and several other advancements in the domain of Explainable AI (XAI) now allow us to understand how the model arrives at specific credit decisions. This not only enhances trust in the credit scoring process but also addresses the common critique that AI models are "black boxes."



Gen AI has the potential to revolutionize credit scoring and assessment with its ability to create synthetic data and understand intricate patterns, offering a more nuanced, adaptive, and predictive approach. Gen AI's capability to synthesize diverse data sets addresses one of the key limitations of traditional credit scoring—the reliance on historical credit data. By creating synthetic data that mirrors real-world financial behaviors, gen AI models enable a more inclusive assessment of creditworthiness.

Adaptability plays a crucial role in navigating the dynamic nature of economic conditions and changing consumer behaviors. Unlike traditional models that struggle to adjust to unforeseen disruptions, gen AI's ability to continuously learn and adapt ensures that credit scoring remains effective in real-time, offering a more resilient and responsive tool for assessing credit risk.

In addition to its predictive prowess, gen AI can contribute to transparency and interpretability in credit scoring. Models can generate explanations for their decisions, providing clearer insights into credit assessments, and enhancing trust among consumers, regulators, and financial institutions.

Enhance credit applications and scoring with MongoDB in the following key aspects of the process:

Simplify data capture and processing:

streamlining this process by eliminating redundant requests and leveraging existing data could significantly enhance the user experience. Leveraging JSON for online credit application forms simplifies the data capture process and also the performance in data processing. JSON's structured data representation proves highly conducive for organizing the multifaceted information within credit applications, encompassing personal, financial, and employment details.

Enhance credit scoring with AI:

Leveraging [MongoDB's developer data platform](#) — an integrated suite of data services centered around a cloud database — we can create a comprehensive customer/user banking profile by combining relevant data points.

Below is an architectural diagram of the data processing pipeline for the predicting probability of delinquency and credit scoring. The data pipeline for credit scoring a customer involves the following steps: data collection, data processing, risk profile generation, model development, data transformation, and decision collection.

Explain the credit application

declination: when it comes to credit application declination, understanding the reasons behind it is crucial. Let's explore how MongoDB and [large language models](#) (LLMs) can shed light on XGBoost model predictions (the model used in this [tutorial](#)). The risk profiling ML pipeline employed provides a probability score that defines the risk associated with the profile for product recommendation. This message is communicated back to the user in a templated manner where only the final status of the application is communicated to the end user.

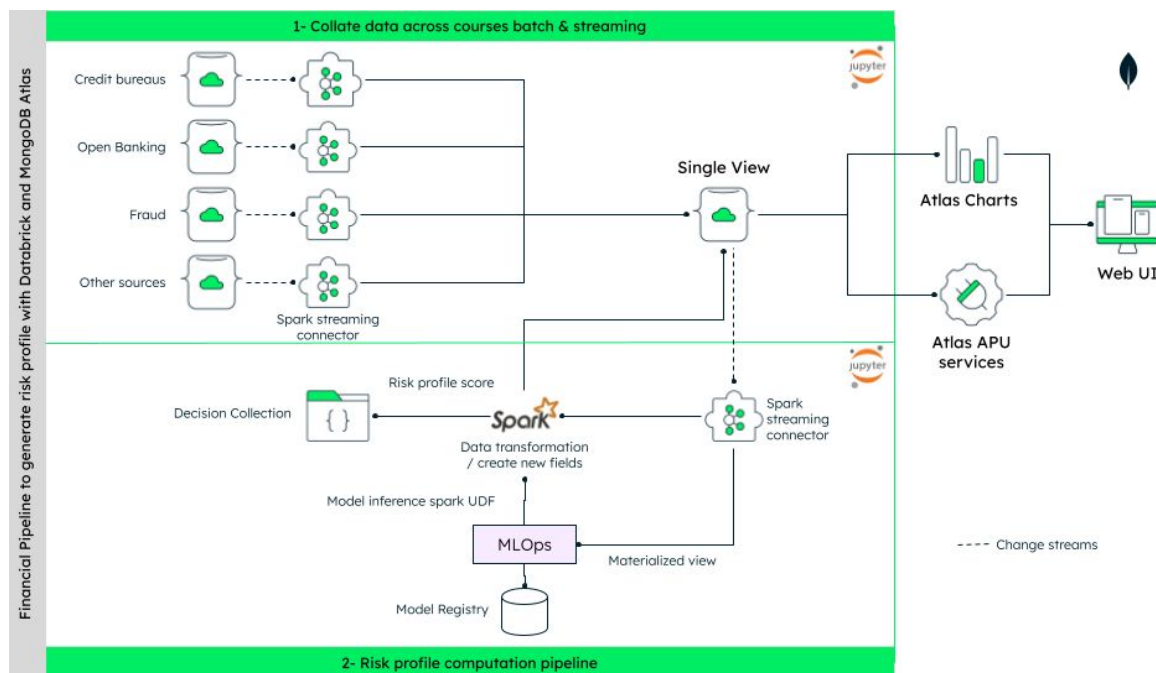


Figure 29: Architecture of the data processing pipeline for the predicting probability of delinquency and credit scoring

Below you will see the architecture explaining credit scoring using an LLM.

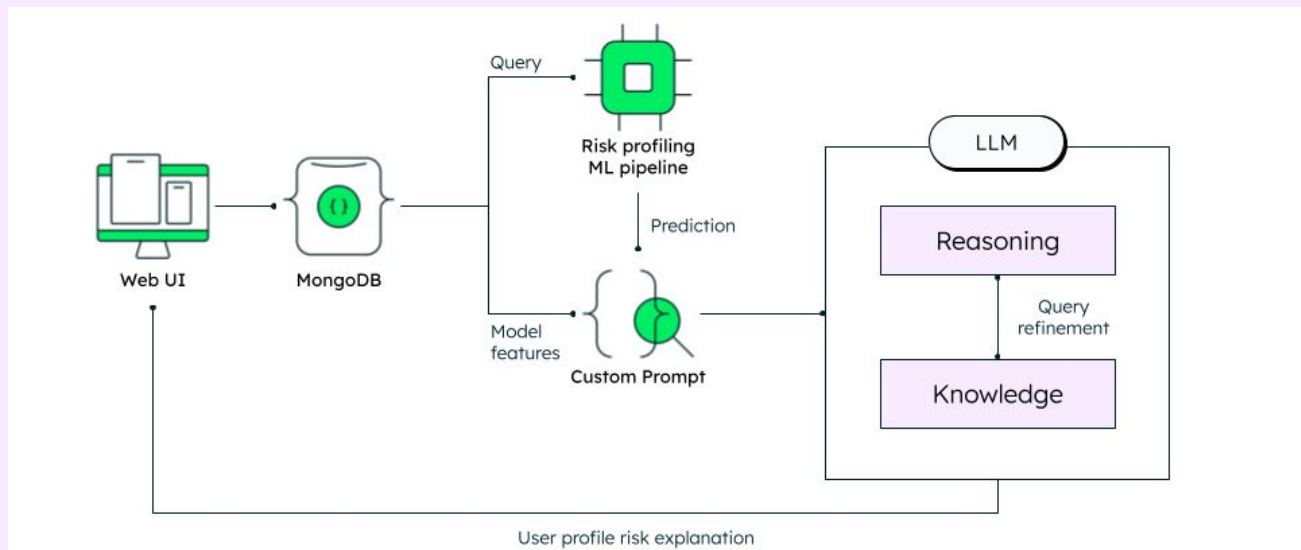


Figure 30: Architecture of the credit application declination

Recommend alternative credit products:

If the credit product applicant is declined, the credit institution should still try to cross-sell to the customer with a relevant product that meets their needs as they are

already engaged in the process. Below is an example of a data architecture that is used to achieve this.

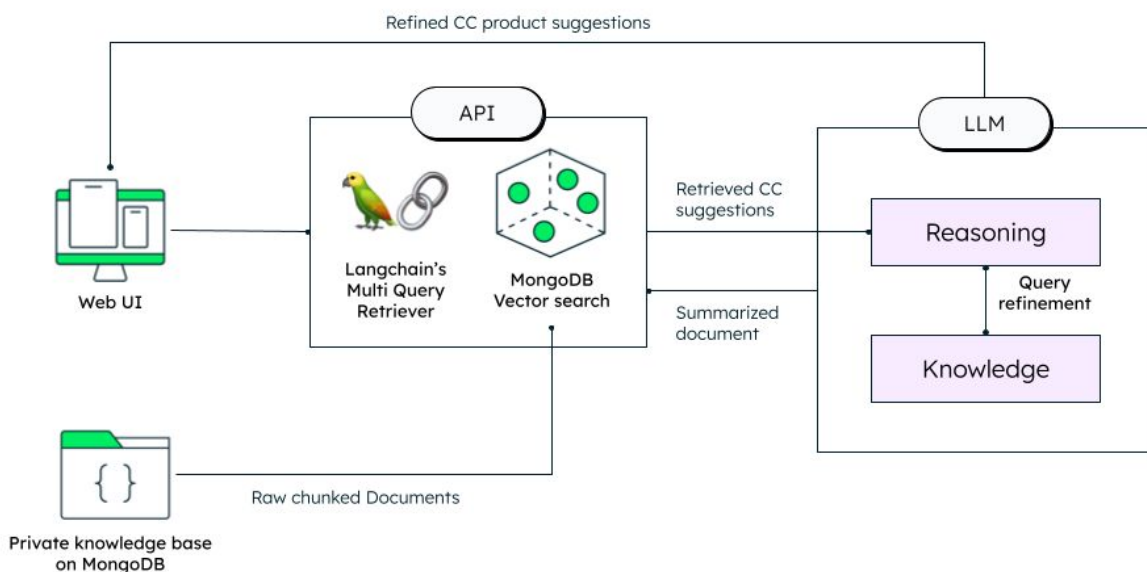


Figure 31: Architecture of the alternative credit product recommendations

Learn how with this credit card Credit card application with generative AI tutorial

Assessing Business Loan Risks

Business loans are a cornerstone of banking operations, providing significant benefits to both financial institutions and the broader economy. In 2023, the value of commercial and industrial loans at commercial banks in the United States reached nearly 2.8 trillion U.S. dollars. However, these loans involve significant risks, including credit and business risks.

Importance of Business Plans

A comprehensive business plan is essential in securing business loans, providing detailed insights into a borrower's strategy, market position, operational plans, and financial forecasts. However, the complexity and length of these plans pose challenges for loan officers, including time constraints and the need for specialized knowledge to interpret key metrics and risk factors accurately.

Generative AI in Risk Assessment

Generative AI can streamline the analysis of business plans by extracting essential information and highlighting key risks. This includes using gen AI-powered chatbots that simulate discussions with loan officers, providing real-time risk assessments based on the input data.

For instance, ChatGPT-4.0 can be prompted to perform a flood risk assessment. However, due to insufficient flood data, it first indicated the location was in a low flood risk zone but later claimed it had significant flooding. This inconsistency can confuse users and lead to incorrect risk assessments.

The screenshot displays a chat window with a light blue header. The user's prompt asks for business risks, profitability projections, and repayment ability for a \$100,000 loan for a bakery at 1633 Broadway, New York, NY 10019, USA. The AI's response is structured with a title 'Starting a bakery...' followed by an overview and a section '1. Business Risks'. Under 'Market Competition', it lists 'High Competition' and mentions that the full risk assessment is more detailed. A follow-up prompt asks for a flood risk assessment. The AI responds with 'Flood Risk' and '1. Location Analysis:', stating that 1633 Broadway, NY, is generally not in a high flood risk zone but that FEMA flood maps and local history should be reviewed. A second prompt asks if there has been any flood within 5km. The AI responds with '2. Flood Risk:', mentioning recent flood events like Tropical Storm Henri and Hurricane Ida.

What are the business risks, profitability projections, and repayment ability of a \$100,000 loan for starting a bakery at 1633 Broadway, New York, NY 10019, USA?

Starting a bakery at 1633 Broadway, New York, NY, a prime location in Midtown Manhattan, involves careful consideration of several factors, including business risks, profitability projections, and the ability to repay a \$100,000 loan. Here's an overview of these aspects:

1. Business Risks

Market Competition

- **High Competition:** Midtown Manhattan is a bustling area with numerous established bakeries, cafes, and food outlets. New entrants face significant competition from well-known brands

The full risk assessment is more detailed

assume you are a loan officer analyzing the risk of this business. What is the business risk flood risk, profitability projection and repayability of a 100000USD loan for starting a bakery at 1633 Broadway, New York, NY 10019, USA?

Flood Risk

1. **Location Analysis:**

- 1633 Broadway, NY, is generally not in a high flood risk zone, but it's essential to review FEMA flood maps and local flood history.

has there been any flood within 5km of 1633 Broadway, New York, NY 10019, USA?

Searched 4 sites

2. **Flood Risk:**

- **Recent Flood Events:** The area has experienced significant flooding in recent years due to storms such as Tropical Storm Henri and Hurricane Ida, both of which brought heavy rainfall

Figure 32: Examples of how ChatGPT-4.0 could respond when asked to assess the risk of a business loan and how it responded to flood risk as a factor

Retrieval-augmented generation (RAG) and MongoDB


While gen AI offers significant potential, its effectiveness in loan officer workflows is hindered by the need for constant prompting and context augmentation. This manual process is time-consuming and often requires specialized skills. To overcome these challenges and enhance data accuracy, we propose a retrieval-augmented generation (RAG) approach utilizing MongoDB's multimodal capabilities. By incorporating geospatial data and other external datasets, RAG can provide a more comprehensive risk analysis. For instance, MongoDB's geospatial queries can identify historical flood occurrences near a business location, enriching AI-generated insights with crucial context.

Demonstration and Results

Let's examine several figures from a generative AI RAG demo built using LLama 3 on the Fireworks.AI platform, with MongoDB as an operational data store for geospatial queries. In this demo, users are asked for a business location using a map (where they can either choose one of the defaults from a dropdown list or use a geolocation pin to pick a location on the map) and a short business purpose and description (to emulate the idea of a business plan).

Please indicate the business location in the USA for your real estate by clicking on the map, or entering the address on the search bar.

Q 88 Kearny St Suite 500, San Francisco, CA 94108, United States



Leaflet | Selected Coordinates: 37.7886378, -122.4033442

Please provide a brief description of your loan purpose and business plan.

Vibrance Studios is a unique, membership-based art studio offering a variety of creative classes and workshops for adults and children. Our studio will provide a welcoming space for individuals to explore their artistic side, relax, and socialize. We'll offer classes in painting, pottery, sculpture, and more, taught by experienced local artists.

We're seeking a \$400,000 loan to cover startup costs. The funds will be allocated as follows: \$200,000 for studio design and build-out, \$100,000 for initial marketing and advertising efforts, and \$100,000 for working capital to cover instructor fees, supplies, and rent until we reach profitability.

One potential risk is competition from established art studios in the area, but we believe our unique membership model and diverse class offerings will set us apart. Additionally, we'll focus on building a strong online presence and community to attract and retain members.

Clear Example Submit

Figure 33: Choosing a location on the map and writing a brief plan description

Upon submission, the AI performs a RAG analysis. Using prompt engineering, it delivers a simplified business analysis considering the location and flood risk data from external sources. [McKinsey's survey](#) found that 65% of respondents report their organizations are regularly using generative AI, including RAG, in at least one business function. This demonstrates the increasing adoption and integration of RAG within business processes.

Assessor's response

As a risk assessor for a business loan company, I have analyzed the business description, address, and flood data for Vibrance Studios. Here is my assessment:

Business Description and Cost/Revenue Projections:

Vibrance Studios is a unique, membership-based art studio offering a variety of creative classes and workshops for adults and children. The business is seeking \$400,000 to cover startup costs, with \$200,000 allocated for studio design and build-out, \$100,000 for initial marketing and advertising efforts, and \$100,000 for working capital to cover instructor fees, supplies, and rent until profitability. Based on industry benchmarks, I estimate the average monthly revenue to be around \$40,000, with a projected annual revenue of \$480,000. The break-even point is estimated to be around 8-10 months after launch.

Local Competition and Demographics:

The location of Vibrance Studios in San Francisco, CA, is a densely populated area with a high demand for creative activities. The demographics suggest a diverse population with a strong interest in art and culture. However, there is moderate competition from established art studios in the area, which may impact market share.

Flood Risk Assessment:

The flood data indicates a high risk of flooding at the proposed location. There were 7 floods less than 5km away last year, and a total of 44 floods less than 5km away over the past 5 years. This suggests a recent and frequent history of flooding in the area. Although the number of floods over 5km away is higher, the proximity and recency of the floods less than 5km away are a significant concern. The flood risk is further exacerbated by the fact that the business is located in a flood-prone area, as indicated by the University of Oklahoma's crowdsourced database, Dartmouth Flood Observatory, and other sources.

Risk Analysis:

The high risk of flooding poses a significant threat to the business, particularly in terms of property damage and disruption to operations. The lack of flood insurance coverage was not mentioned in the business description, which is a concern. The business may need to consider investing in flood mitigation measures, such as flood-resistant construction and emergency preparedness plans. Overall, while Vibrance Studios has a unique business concept and a growing market, the high risk of flooding at the proposed location is a significant concern. The business will need to carefully consider and mitigate this risk to ensure long-term success.

Figure 34: Loan risk response using RAG

Clicking the "Pin" icon next to the location bar reveals sample flood locations near the selected business site. The pins mark flood locations, and a blue circle shows a 5km radius where flood data is queried.

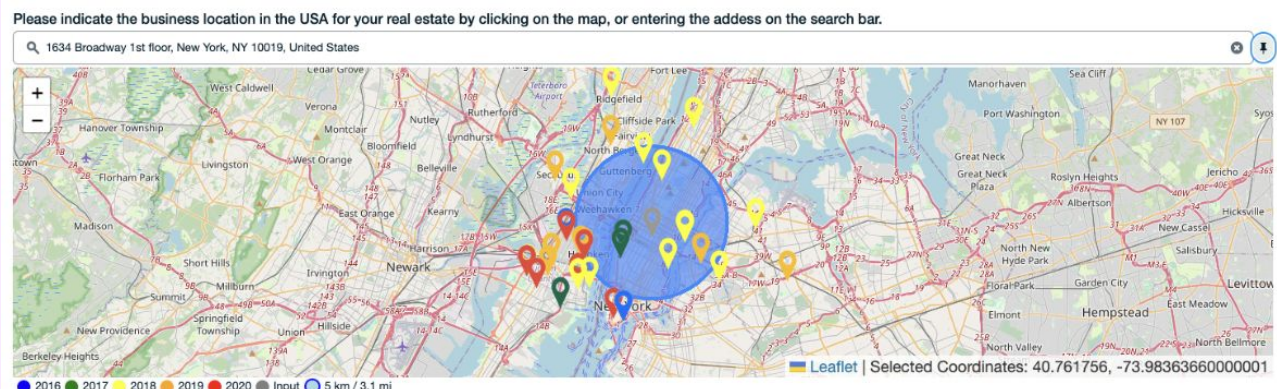


Figure 35: Flood locations displayed with pins

RAG and multimodal data: a robust approach to loan risk assessment

Integrating gen AI with a multimodal developer data platform like MongoDB Atlas offers a robust solution for detailed and accurate business loan risk assessments. The iterative nature of the RAG process ensures continuous learning and improvement, minimizing risks like data hallucinations. As AI models become increasingly multimodal, leveraging comprehensive data platforms will be crucial for maximizing their potential.

Solution: Better Business Loans with MongoDB and Generative AI.

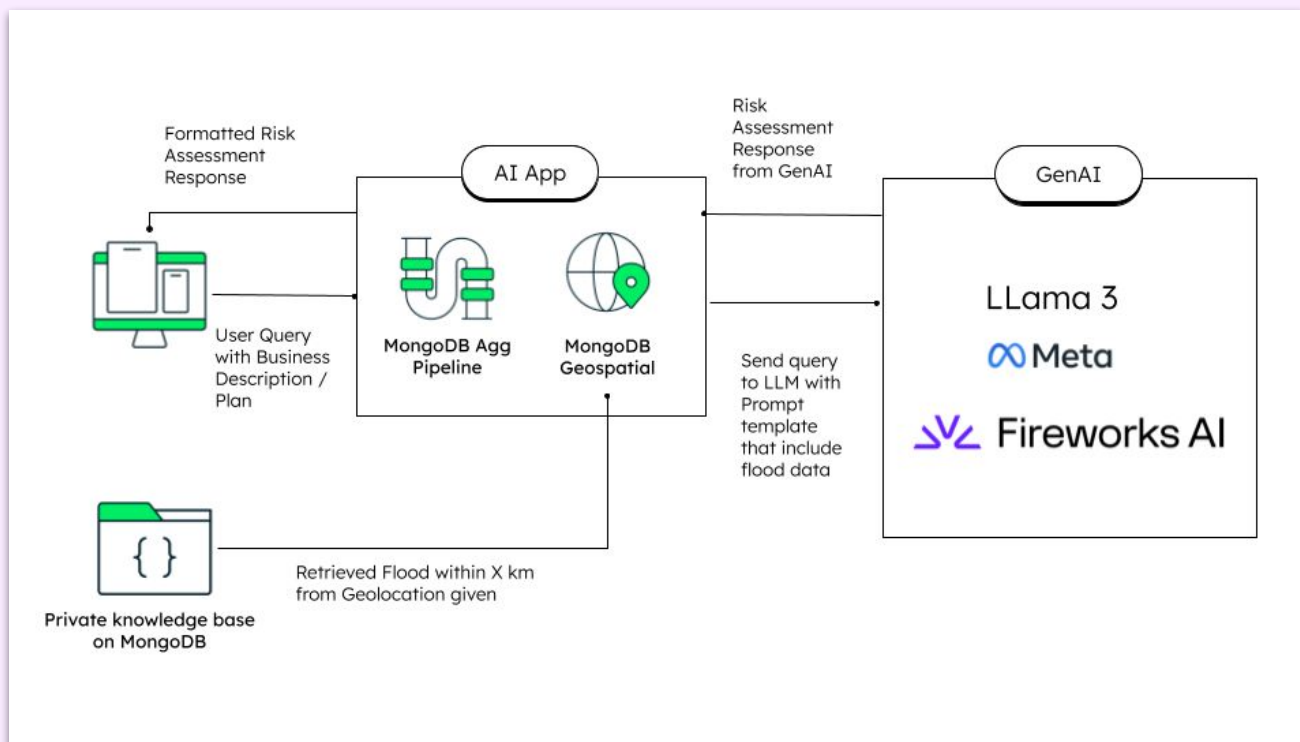


Figure 36: RAG data flow architecture diagram

Emerging use cases for AI in payments

A lack of developer capacity is one of the biggest challenges for banks when delivering payment product innovation. Banks believe the product enhancements they could not deliver in the past two years due to resource constraints would have supported a [5.3% growth in payments revenues](#). With this in mind and the revolutionary transformation with the integration of AI, it is imperative to consider how to free up developer resources to make the most of these opportunities. There

are several areas in which banks can apply AI to unlock new revenue streams and efficiency gains. The image below provides a high-level view on eight of the principal themes and areas. This is not an exhaustive view but does demonstrate the depth and breadth of current opportunities. In each example, there are already banks that have begun to bring services or enhancements to market using AI technologies or are otherwise experimenting with the technology.

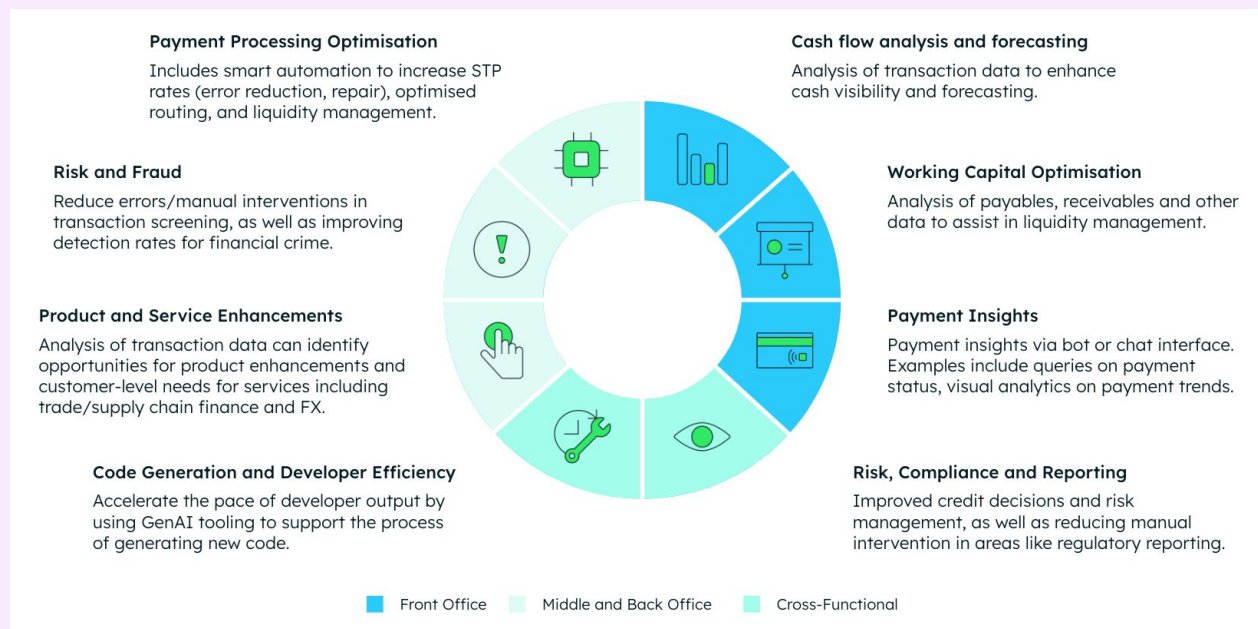


Figure 37: Benefits of payment processing optimization

Below are some areas [according to tech analyst Celent in which AI can be leveraged](#) across financial institutions in front, middle and back office operations.

Front office for new and richer insights:
AI has huge potential to support

value-added services making the most of payments data. Analysis of payables, receivables, and balance information in aggregate can be used to deliver a range of different insights and service improvements in real-time. Below are three areas where

financial institutions can use AI to enhance their customer-facing services:

- **Cash flow analysis and forecasting:** Using the information in payment messages to provide granular and real-time insights into the cash position for a client.
- **Working capital optimization:** Using AI in conjunction with payments data can support a range of use cases around liquidity and management with insights and recommendations.
- **Payment insights:** Gen AI can support the delivery of payment and account insights to customers.

Middle and back office for workflow and process improvements:

AI and advanced analytics are already commonplace across large areas of middle and back-office operations for payment processing, but many potential opportunities for further enhancement still exist:

- **Payment processing optimization:** One use case for machine learning is to further automate the payments repair process. AI technologies to improve transaction routing is another important opportunity. A third would be in message translation or conversion between older standards and ISO formats, e.g. to translate unstructured name and address in ISO 8583 messages to structured ISO 20022 messages.
- **Risk and fraud:** As with the use case for payments repair, applying AI technologies in transaction screening can reduce the need for manual intervention.

- **Product and service enhancements:** Personalizing services is a huge area AI can help.

Cross-functional capabilities to increasing agility and efficiency

AI technologies can also be used to support several cross-functional use cases within a bank:

- **Using gen AI to support code generation:** This is a huge opportunity for institutions to increase developer efficiency.
- **Risk, compliance, and reporting:** Technologies such as machine learning, NLP, computer vision, and gen AI can all be deployed in various combinations to support processes such as risk management and regulatory reporting.

Dive into Celent's report to learn more about the use cases above and understand how you can harnessing the benefits of AI in payments.



Regulatory Compliance and Code Change Assistance

The regulatory landscape for banks has grown increasingly complex, demanding significant resources for the implementation of numerous regulations. Traditionally, adapting to new regulations has required the manual translation of legal text into code, provisioning of data, and thorough quality control—a process that is both costly and time-consuming, often leading to incomplete or insufficient compliance. For instance, to comply with the Basel III, Dodd-Frank, PSD2, or various other regulations, developers must undertake extensive coding changes to accommodate the requirements laid out in thousands of pages of documentation.

AI has the capacity to revolutionize compliance by automating the translation of regulatory texts into actionable data requirements and validating compliance through intelligent analysis. This approach is not without its challenges, as AI-based systems may produce non-deterministic outcomes and unexpected errors. However, the ability to rapidly adapt to new regulations and provide detailed records of compliance processes can significantly enhance regulatory adherence.

Recent research highlights that generative AI not only identifies regulatory obligations and gaps but also provides accurate code recommendations. For example, generative AI can help a bank ingest all applicable regulatory requirements and internal controls, quickly alerting it to any discrepancies. By harnessing its comprehensive and sophisticated capabilities, generative AI can significantly reduce manual effort and minimize errors, all while strictly adhering to factual data.

It then determines an organization's obligations and evaluates its compliance. If obligations are unmet, AI maps them to specific controls or generates new ones. As AI models continue to train, their accuracy

in identifying gaps and interpreting new regulatory requirements improves. Eventually, generative AI may even write new controls, driving greater efficiency.

This seamless process is facilitated by using a multimodal database like MongoDB, which addresses the common issue of documents being scattered and lacking proper metadata and documentation. MongoDB offers several key differentiating features:

- **Converged operational and AI data storage:** MongoDB allows organizations to keep all their operational, unstructured, and vector data in one place, eliminating data silos and making it easier to access and manage.

- **Efficient data handling:** MongoDB simplifies working with data at scale. Centralizing data reduces manual effort, leading to quicker and more accurate compliance processes.
- **Enhanced search and generative AI capabilities:** MongoDB facilitates easy searching through documents and metadata, crucial for quickly identifying regulatory requirements and checking them against existing controls. Utilizing RAG capabilities with MongoDB provides better context to LLMs for generating code suggestions.

Below is an example of a conceptual workflow using MongoDB for code generation:

1. Regulatory requirements: Gather relevant regulatory requirements, determine the specific obligations from these regulations, and condense the regulatory text into concise summaries.

2. Translate to controls: Convert these summaries into actionable controls. *In this context, **control** is a specific measure often implemented through code that ensures compliance with regulatory requirements by automating necessary actions, such as enforcing multi-factor authentication (MFA) for certain transactions.*

3. Convert into executable code:

- Check the controls against the organization's existing internal controls, identify gaps, and create new controls where needed.
- Convert controls into executable code.

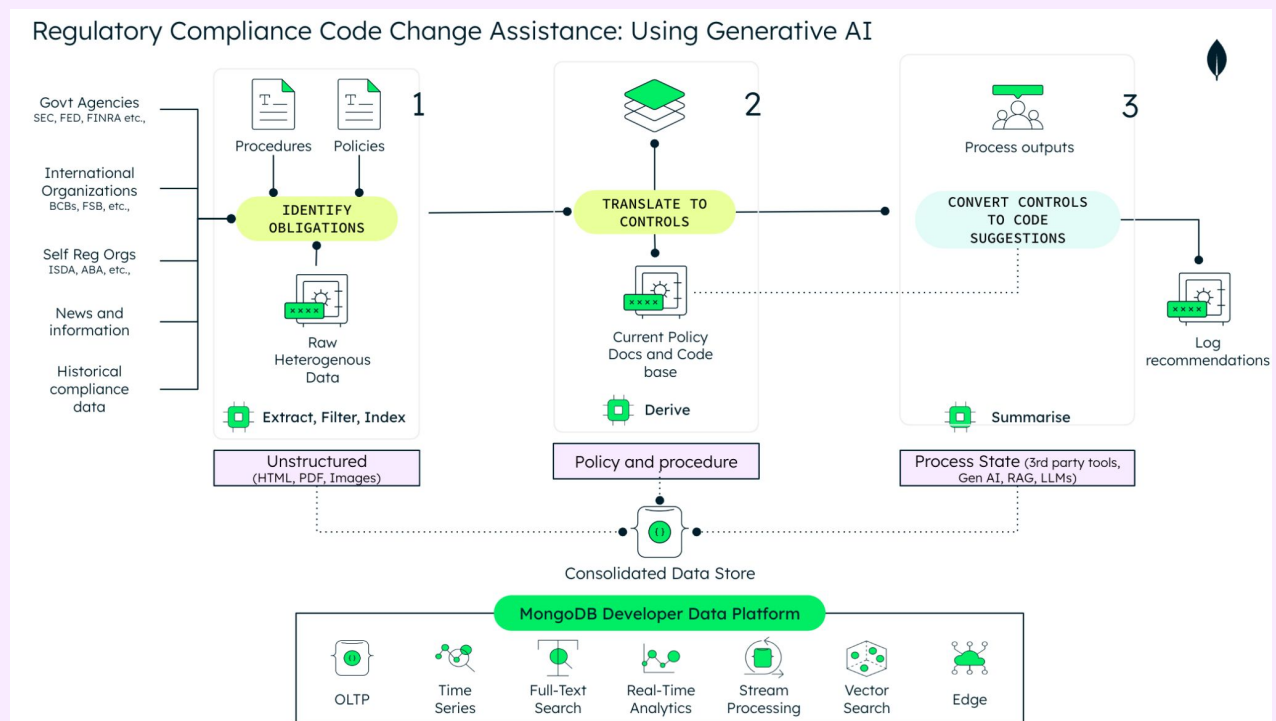


Figure 38: Automated compliance code adaptation with AI

In the financial services section, we emphasized the importance of regulatory compliance by referencing several key regulations such as AML, Basel III, PCI, and GDPR. To demonstrate the practical application of code generation, we used PSD2 as an example and engaged ChatGPT to see if this approach is feasible. Below, you will find the ChatGPT prompts and responses.

Step 1: Real PSD2 clause

Banks must apply strong customer authentication (SCA) for electronic payments initiated by the payer, unless an exemption applies.

Step 2: Control

Ensure that any electronic payment initiated by the payer requires multi-factor authentication (MFA), unless an exemption applies.

Step 3: Code suggestion

Using MongoDB and generative AI, the process might look like this:

1. **Identify payment transactions:** Store all electronic payment transactions in MongoDB.
2. **Check transaction type:** Create a control to check if the transaction is initiated by the payer.
3. **Implement MFA:** Generate code to enforce multi-factor authentication (MFA) for applicable transactions.

The full example is more detailed

please provide an example of converting PSD2 regulation to code



Certainly! Let's walk through an example of converting a specific PSD2 regulatory requirement into code. We'll focus on the Strong Customer Authentication (SCA) requirement, which mandates the use of multi-factor authentication (MFA) for online payments.

Figure 39: ChatGPT 4.0 Example of PSD2 regulation converted into code

Below is a simple example of what the generated code for MFA might look like.

```
from pymongo import MongoClient
from datetime import datetime

client = MongoClient("mongodb://localhost:27017/")
db = client["bank"]
transactions = db["transactions"]

# Function to enforce Strong Customer Authentication
def enforce_sca(transaction):
    if transaction["payer_initiated"] and not transaction["exemption"]:
        # Example MFA enforcement (pseudo-code)
        send_mfa_request(transaction["user_id"])
        return "MFA required"
    return "MFA not required"

# Insert a new transaction and enforce SCA
new_transaction = {
    "transaction_id": "78910",
```

Figure 40: MongoDB setup for multi-factor authentication

```
    "user_id": "user123",
    "amount": 50, # Amount in euros
    "payer_initiated": True,
    "exemption": False,
    "transaction_time": datetime.utcnow()
}

# Insert transaction into the database
transactions.insert_one(new_transaction)

# Enforce SCA for the new transaction
result = enforce_sca(new_transaction)
print(result)

# Example function to send MFA request (pseudo-code)
def send_mfa_request(user_id):
    print(f"Sending MFA request to user {user_id}")
```

Figure 41: Transaction handling and strong customer authentication (SCA) enforcement

Other Notable Use Cases



Below are a few more examples of where AI can be leveraged in financial operations:

Risk modeling: AI can create synthetic scenarios and data that can be used to stress test financial systems and models

Algorithmic trading: AI algorithms can analyze public market data and execute trades at speed, optimizing trading strategies.

Generating financial reports: AI can analyze financial data including transactions, invoices, and account statements to automate the process of generating reports. By using AI and ML techniques, relevant information can be extracted where required.

These examples highlight several avenues for integrating AI within financial institutions. Embracing AI in financial applications promises enhanced risk management, operational efficiency, and superior customer experiences.

Therefore, it is essential for financial institutions to grasp the profound technological implications, scale, and intricacies associated with AI, particularly in crafting a generative AI strategy. Adopting a strategic and holistic approach that addresses technological, data, ethical, and organizational dimensions is imperative for navigating this transformative landscape effectively.



Bendigo and Adelaide Bank uses gen AI and MongoDB to modernize applications

Bendigo and Adelaide Bank is a leading Australian financial institution with over two million customers and 7,000 employees. By prioritizing digital transformation, the bank successfully migrated 32 banking applications to AWS in just 30 days, leveraging MongoDB to accelerate development and deliver enhanced services to its customer base.

Contact Information



Wei You Pan

Global Director, Financial Services Industry
weiyou.pan@mongodb.com



Luis Pazmino Diaz

Principal, Financial Services Industry
luis.pazmino@mongodb.com



Peyman Parsi

Senior Principal, Financial Services Industry
peyman.parsi@mongodb.com

FOR MORE INFORMATION AND RESOURCES

Visit MongoDB for Financial Services

