

HOW FINANCIAL SERVICES LEADERS CAN ADOPT A SCALABLE GENERATIVE AI STRATEGY

Key considerations for financial institutions to fully embrace gen AI

TABLE OF CONTENTS

- 3 Introduction
- 4 Generative AI Use Cases for Financial Services Firms
- 6 Challenges Facing Generative AI in the Financial Services Sector
- 8 Financial Services Firms Require a Robust Data Strategy
- **9** An Essential Tool for the Future of Financial Services

INTRODUCTION

No other private sector has embraced advanced technology to the degree that financial services has. Machine learning (ML) and AI form the backbone of cybersecurity and fraud detection operations for this **estimated \$28 trillion global industry**.

High-performance computing, combined with AI models, is used to efficiently extract and predict trading patterns from billions of transactions each day, or parse thousands of pages of unstructured data in security filings, policies and other documents to efficiently scale the set of products financial services companies can offer their customers.

In the world of predictive AI, also known as machine learning, building the models used for prediction or automation once required advanced coding and math skills to develop and use. But now in the age of generative AI (gen AI), using pre-trained models from companies such as Mistral AI, Meta, OpenAI, NVIDIA, Microsoft Research and many more, anyone can harness the benefits of AI using natural language. It has a democratizing effect, putting this technology in the hands of every business user who may be less aware of the reputational and regulatory risks of using proprietary or sensitive information with one of these models, especially when hosted outside of their organization's governed environments.

That can make financial services stakeholders in such a highly regulated industry understandably nervous, especially since the adoption of these tools is already well underway. A **KPMG survey** of 300 financial service executives in June 2023 found that roughly half had implemented gen AI solutions, and another 35% planned to do so within six months. More than a third of companies surveyed by KMPG are relying on off-the-shelf products such as ChatGPT, while 53% are developing home-grown solutions, a signal of both innovation and the need to address the serious privacy risks of using proprietary data in externally hosted models.

During a period of heightened competition, gen AI can help financial services firms deliver new, innovative solutions and meet escalating consumer expectations without forgetting the need to maintain regulatory compliance. Gen AI is not just a tool but rather a way to give every employee and customer a personal assistant. Gen AI can drive new revenue streams, help companies operate more efficiently and drive down costs via advanced automation and augmentation.

The key to realizing the potential of gen AI is by first having a solid data strategy in place. One with strong governance practices that focuses on eliminating silos and preserving privacy throughout the data lifecycle – from when it is stored to when it is used with an LLM, and results from that data are surfaced via an application such as a chatbot.

Let's delve into the adoption challenges facing financial services companies, how jobs and workflows may be impacted, and how you can navigate the evolving AI landscape to remain at the forefront of innovation and competitiveness.

How Gen Al Powers Conversations with Data

Gen AI applications are built on top of large language models (LLMs). LLMs are deep neural networks trained to understand the meaning and structure of human language by exposing them to enormous volumes of text. This allows any user to have back-and-forth conversations with the model.

The ability to "converse" with gen Al models in a natural way eliminates the need for special syntax or technical expertise, removing many barriers to Al adoption. Gen Al's ability to interact conversationally and summarize content drawn from a deep well of data allows it to be applied across a wide range of business use cases.

Foundation LLMs are often trained primarily with internet data. Because these widely-available foundation models have not been trained on proprietary data, they may "hallucinate" (i.e., provide incorrect responses) when providing outputs in an enterprise use case requiring specific context. To deliver true value in the financial services, it is ideal for foundation LLM to be further trained leveraging proprietary data sets.

GENERATIVE AI USE CASES FOR FINANCIAL SERVICES FIRMS

Financial services firms have always been among the first to embrace AI, but its use has been largely limited to data science specialists and a small subset of business analysts. Gen AI opens up access to this powerful tool to virtually everyone within the organization. This enables a range of new use cases across three broad categories: personalizing the customer experience, enabling more efficient market intelligence and portfolio management, and streamlining risk and compliance processes.



ENHANCING THE CUSTOMER EXPERIENCE

Every organization wants to use data to better understand its customers. Delivering a more personalized experience can be a key differentiator in a commodity market such as banking or insurance. Gen AI helps to accelerate workflows in the following areas:

• Intelligent chatbots and virtual assistants:

Automated service agents can take advantage of advanced natural language processing (NLP) to hold human-like conversations with customers, offering a more personalized level of service while saving costs on call center operations.

- Automated advisory services: As with chatbots, gen Al-based robo-advisors can combine algorithmdriven financial advice with the ability to answer questions from investors and banking customers. They can also assist with tailoring investment strategies to a customer's financial objectives and risk-tolerance profile.
- Smoother claims processing: Gen AI's ability to understand huge volumes of unstructured data, such as PDFs of repair bills and photographs of damaged vehicles, can help streamline and automate the claims process, resulting in higher customer satisfaction. In addition, the ability to identify potential high-cost claims earlier in the process can help insurers allocate resources more efficiently.

IMPROVING MARKET INTELLIGENCE AND PORTFOLIO MANAGEMENT

The ability to assess vast volumes of data and make quick decisions is the key to staying ahead of market trends. What was once the sole territory of trade desk specialists is now accessible to anyone with access to a foundation LLM trained on the relevant financial data. Use of LLMs may provide benefits such as:

- More efficient analysis of market sentiment: Firms can use LLMs to streamline the analysis and summarization of unstructured and semistructured text from SEC filings, regulatory reports, traditional mass media publishers and social media sources. Then with human validation they can use that data to infer market sentiment toward specific topics and the impact on investment strategies.
- Automated portfolio risk management: Sophisticated AI simulations can ingest market sentiment analysis, exchange rate predictions, interest rate scenarios, historical weather patterns and many other data sources, then use them to validate and test models for predicting portfolio risk.
- Better portfolio optimization: Gen AI systems can simulate a massive number of market conditions, product and price scenarios, interest rate models and other variables. These simulations can be used to develop adaptive trading strategies that learn over time to optimize trading and improve portfolio performance.

STREAMLINING RISK AND COMPLIANCE

For financial services firms, regulatory compliance is job number one. It's a huge area of investment for banks, insurers and trading firms, and more and more of those dollars are being spent on AI solutions to combat increasingly sophisticated attempts at financial crime. Here, gen AI can help in the following ways:

- More comprehensive stress testing: All financial institutions are now required to demonstrate they can withstand a severe economic crisis. By leveraging gen Al, risk and compliance teams can simulate a wide range of adverse market conditions based on historical events, current market conditions and potential future risks. Gen Al can also be used to generate synthetic data to test the validity of the financial models used in stress testing.
- Proactive risk management: Gen AI can help risk teams create models that assess and forecast a bank's exposure across fluctuations in interest rates, credit, liquidity and the possibility of default. It may also be used to identify the warning signs of corporate bankruptcies, anticipate likely mergers and acquisitions and predict other events that impact risk assessment.

• Anticipate and prevent unknown fraud vectors: Gen AI can be used to anticipate new patterns of fraudulent activity that don't yet exist. By drawing up variations in synthetic transactions, security teams can identify previously unknown fraud patterns, which can then be used to update detection and prevention models.

In addition to these benefits, gen Al tools can also be used for standard sales and marketing functions, such as creating highly personalized emails, creating first drafts of marketing collateral and other materials, and identifying product suggestions and deals tailored to each customer.

Generative Al captured the industry's attention last year. But now as organizations evolve from hype to implementation, many face a reality that also comes with many questions. But one certainty remains: a robust and scalable AI strategy hinges on an end-to-end enterprise data strategy. If done successfully, organizations can harness the full power of data and drive exponential productivity gains."

-Rinesh Patel, Global Industry GTM Lead for Financial Services at Snowflake

CHALLENGES FACING GENERATIVE AI IN THE FINANCIAL SERVICES SECTOR

In a highly regulated environment such as financial services, gen AI may pose additional challenges around data privacy, security and compliance. Organizations may also grapple with difficulties implementing the technology effectively and adapting their business to optimize gen AI's effectiveness.



DATA CHALLENGES

- Securing data during training: Organizations planning to use their proprietary data to finetune an off-the-shelf foundation model could risk exposing that data to other users of that model, depending on the security of their training environment. The training data may include personally identifiable information (PII) and/or intellectual property.
- Securing data during use: Because some gen Al systems may learn based on the prompts entered by their users, individuals could inadvertently share personal or proprietary data while using them. That data could then be subsequently exposed to people outside the organization.
- Managing regulatory compliance: Beyond addressing issues of data privacy and security, financial services companies will need to ensure that any outputs generated by gen AI are compliant and consistent with industry expectations and government regulations. Even more critical is ensuring that advice generated by robo-advisors complies with relevant laws and regulations. Without such compliance, customer trust itself is at risk.

TECHNOLOGY CHALLENGES

- Managing technical complexity: Very few enterprises possess the resources to develop their own LLMs. This is why most organizations will fine-tune or run customizations on top of pretrained foundation models (such as GPT4, BERT, LLaMA or Mistral AI) using their own data. But even then, data science expertise will be required. The foundation model must be fine-tuned specifically to the organization's data, requiring expertise in model governance and gen AI development.
- Meeting high resource demands: Developing a foundation model requires enormous amounts of structured and unstructured data, which must be stored and managed within a cohesive environment such as a data lake. The training process will also require significant amounts of computing power, typically massively parallel banks of GPUs.
- Minimizing data movement across applications: Fielding requests for AI-powered apps from every business unit puts pressure on internal developer teams to create siloed copies of data, and opens data risks when teams onboard third-party apps with LLMs that access their own data. A modern cloud-based data platform, where the data and the LLM remains within a governed boundary, is a more secure approach to scaling adoption of gen AI applications.

ORGANIZATIONAL CHALLENGES

• Realigning your workforce around AI:

Whether gen AI systems will increase employee productivity, replace some workers or simply automate many routine tasks, the ways in which it may free up employees' time to perform more strategic work will likely vary from one organization or department to the next. Either way, organizations must consider how to manage what is likely to be a fundamental shift in how their labor force functions.

 Preventing Al hallucinations and other inaccuracies: Gen Al systems may occasionally produce content that sounds true but isn't – a phenomenon known as a hallucination. This may pose compliance risks and may even open the door to legal implications. Financial services companies must take special care to fact check any materials produced by gen Al systems to ensure that hallucinations do not slip through the cracks.

• Mitigating intentional misuse and abuse:

Gen AI models can be maliciously used to generate deliberate misinformation, posing a compliance risk, or they may be prompted to generate code with intentional "back doors" that make an enterprise vulnerable to cyberattacks. Ensuring proper guardrails around AI use within the enterprise is critical.

- Avoiding algorithmic bias: Concerns about inherent bias have plagued ML models for years, and gen AI is not an exception. Insufficiently diverse or poorly vetted training data can taint the outputs produced by gen AI models. For example, a model trained on historical loan performance data skewed to a particular demographic may produce unfair creditworthiness assessments for individuals who are not members of that demographic. Training data lineage is important to consider before relying on a model to make decisions.
- Training your workforce to use AI the right way: It is likely that all employees will need to adapt their workflows to accommodate gen AI tools. But in addition to teaching workers how to use these tools, education in the ethics of how and when to use AI will likely be required.

FINANCIAL SERVICES FIRMS REQUIRE A ROBUST DATA STRATEGY

The fundamental ingredient in every enterprise AI strategy is data. An AI model is only as good as the data it uses to make predictions. But that data cannot exist in a silo, or multiple silos. It needs to be clean, reliable, accessible and governable.

Before any financial services organization begins to implement gen AI, it first needs to deploy a robust data strategy. Integral to an organization's strategy is a modern data platform that collects all relevant data in one place, secures it, shares it and allows the enterprise to apply uniform and granular governance policies to it. That platform is most likely to be cloudbased, which allows for easier collaboration between internal and external stakeholders, as well as the ability to ingest data from outside sources such as partners, agencies and third-party data vendors.

Here are three key steps you must take to implement a robust data strategy:

1. Break down silos: An August 2023 survey by IDC

found that 90% of enterprise data is unstructured, and half of enterprises maintain that unstructured data in one or more internal silos. Keeping data siloed makes it harder to share, which impedes collaboration and slows the pace of innovation. Collecting data in one place is a necessary first step toward getting it ready for AI. A modern data platform can gather all of the data within a single location and make it accessible for training purposes without compromising its integrity or putting it at risk.

2. Build a flexible, scalable data infrastructure:

Modern IT environments are rarely homogeneous, and this is especially true for financial services firms. Hybrid on-premises and multi-cloud architectures are the norm. Regardless of the diversity of cloud and on-prem systems in the mix, financial services firms need the ability to coalesce huge volumes of first- and third-party data, while also enforcing strict data governance and security policies. A modern data platform is critical for securely sharing and analyzing data across the organization without compromising any sensitive or proprietary information, or running afoul of security mandates.

3. Leverage modern data cloud platforms: Modern cloud-native platforms like Snowflake offer several key advantages over legacy solutions. They make it easier to collaborate with internal and external partners in a secure, policy-driven manner. By applying robust governance policies across all enterprise data, financial services firms can preserve customer privacy, monitor data access on a use-case level, and revoke access when policies are violated. Support for cloud-native technologies such as serverless computing and containers makes it compatible with all major cloud platforms, avoiding lock-in and the onerous ingress and egress fees many vendors charge to move your data between clouds.

SNOWFLAKE CORTEX: THE SMART WAY TO BUILD APPS POWERED WITH GENERATIVE AI

As financial services firms capitalize on the potential of gen AI apps to drive innovation, they'll look to creating their own mini-LLMs using data they can control. **Snowflake Cortex** (currently in public preview on select AWS and Azure regions) is an intelligent, fully managed service that hosts and serves industry-leading AI models, LLMs and vector functions. Cortex allows you to analyze your proprietary data quickly and securely, then build custom AI apps, such as document chatbots, around that data.

Snowflake Cortex is the solution we used to build key applications in the Snowflake Financial Services Data Cloud, including Document Al (which extracts structured data from PDF files), Snowflake Copilot (an LLM-powered coding assistant tailored to your data), and Universal Search (a discovery tool for data clouds) — all in private preview.

Using Cortex's library of serverless SQL and Python functions, financial service firms can quickly build contextually rich applications using the combination of LLMs and vector search functions. Coupled with **Snowpark Container Services**, developers can quickly customize the Al infrastructure components, including any open source LLM or vector database as well as other app components in Snowflake, offering ultimate flexibility from the application UI down to finetuning an LLM.

When using Snowflake, you can also integrate with other trusted managed AI platforms like Azure AI to unlock the full potential of data and AI that meets your specific business needs. The data connector in Azure AI allows you to effortlessly import data into Azure Machine Learning, and with Snowpark External Access you can integrate with API endpoints such as Azure OpenAI Service.

AN ESSENTIAL TOOL FOR THE FUTURE OF FINANCIAL SERVICES

Gen AI facilitates a fundamental shift in enterprise data strategy. Company-wide workflows, software and jobs will be transformed. Every member of the C-suite will need to seriously consider how AI will change the business, and how strong data protection is required to maintain a competitive edge now that everyone has access to the same models. It's also an exciting time for financial services organizations looking for ways to rise above the competition. Gen AI solutions offer multiple paths for incumbents to streamline operations and personalize and add more modern — product offerings. But getting there requires financial institutions to get their data houses in order by adopting a modern, cloud-native data platform focused on ease of use and security.

For more information on how to achieve your Al objectives, visit our **Financial Services Data Cloud**. And if you want to take a deeper dive into gen Al, check out our **Essential Guide to Generative Al**.

A PARTNERSHIP PERFECTLY BUILT FOR THE CLOUD: MICROSOFT AZURE AI

In this new era of AI, it's imperative for financial service organizations to keep pace with warp speed innovation while still working to ensure secure, responsible use of AI tools. Microsoft Azure is an AI-powered cloud platform including more than 200 products and cloud services designed to help you bring new solutions to life - to solve today's challenges and create for the future. Build, run and manage applications on premises, across multiple clouds and at the edge with the tools and frameworks of your choice. The Azure AI portfolio empowers you to innovate faster by providing access to the industry's most cutting-edge models such as OpenAI, Meta, Mistral AI, HuggingFace, Microsoft Research and NVIDIA. Azure AI powers large-scale training of foundation models with the purpose-built AI infrastructure, trusted by leading model builders like OpenAI and Mistral AI. Gain access to the latest open source and frontier models through Azure AI model catalog, and enhance your applications with Azure Al Search's hybrid search capabilities. Simplify model and LLM flow operations for AI developers using built-in tools like prompt flow and Semantic Kernel while embedding responsible AI principles to ensure your AI solutions are secure and ethical from the start. All of these capabilities are seamlessly integrated into a unified generative AI platform, Azure Al Studio, making it easier for you to unlock new potential in AI development.

And when you pair Snowflake and your data with the power of Azure, you can unlock even more value. From easily building gen AI apps — whether you're a coder or not — to customizing Azure OpenAI Service's coding and language AI models with your Snowflake data (while still maintaining high security and governance over it), to creating AI workflows that can improve LLM performance and more, you'll be able to take your team into the future.

EXPLORE AZURE AI



ABOUT SNOWFLAKE

Snowflake enables every organization to mobilize their data with Snowflake's Data Cloud. Customers use the Data Cloud to unite siloed data, discover and securely share data, and execute diverse artificial intelligence (AI) / machine learning (ML) and analytic workloads. Wherever data or users live, Snowflake delivers a single data experience that spans multiple clouds and geographies. Thousands of customers across many industries, including 691 of the 2023 Forbes Global 2000 (G2K) as of January 31, 2024, use the Snowflake Data Cloud to powertheir businesses.

Learn more at **snowflake.com**



ABOUT MICROSOFT AZURE

The Azure cloud platform includes more than 200 products and cloud services designed to help you bring new solutions to life – to solve today's challenges and create the future. Build, run and manage applications across multiple clouds, on premises, and at the edge with the tools and frameworks of your choice. Azure ensures security from the ground up, backed by a team of 8,500+ security experts and a \$20B security investment to support proactive protection that's trusted by enterprises, governments and startups.

Learn more at **microsoft.com**



© 2024 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).