



Capitalizing on generative AI


Part 1 of 3

How banking, financial services and insurance companies can seize the enormous opportunities of gen AI and reduce the equally sizable risk

By Ed Merchant, Craig Weber & Manan Gauba

What's inside

Gen AI: what banking, finance and insurance execs need to know	3
Public pronouncements of gen AI	6
Banking, cards & payments	8
Capital markets and wealth management	9
<hr/>	
The collective impact: potentially huge productivity gains	10
The perilous 11: Key risks may be as big as the rewards	12
Getting prepared without getting overzealous	14
<hr/>	
About the authors	15
<hr/>	
Endnotes	17





Gen AI: what banking, finance and insurance execs need to know



Executive summary

The latest, maybe greatest, advancement in artificial intelligence to date is poised to reshape the operations of every business, but especially financial institutions and insurance companies. At the core of the businesses of financial services and insurance—moving money, providing advice and customer assistance, managing risk and striking deals—is the written and spoken word. Loan officers communicating with risk officers and reviewing written guidelines. Banks and payment contact center reps speaking with, emailing or texting customers. Investment bankers diving into profiles of acquirers and potential targets.

In all these cases, an enormous amount of time is spent manually reviewing text documents. Yet the competitive demands of making rapid decisions undermines their ability to review every possible relevant document, determine what to do and communicate their answers.

This is where generative AI initially plays. Over the last year, many financial institutions and insurance companies have spoken publicly about where they

see its potential. Their excitement is palpable and easy to understand. The technology promises to alter the longstanding ways in which they operate. Generative AI technology uses large language models (LLM), a recent advancement in deep-learning algorithms. It also capitalizes on enormous amounts of economical computing power to search through mountains of digital text, images and numbers in seconds—and deliver well-written answers in seconds.

But the risks of generative AI are just as large as the opportunities to gain serious competitive advantages, and perhaps even larger. One of the biggest is lawsuits alleging questionable decisions, especially about customers (e.g., whom to lend to) and employees (who to hire, promote and fire). The regulatory and legal obligation to explain and justify the rationale for using the output of a complex LLM is giving banks and insurers good reason to be cautious.

Nonetheless, many aren't standing on the sidelines. The most aggressive ones are gauging their risks and experimenting with generative AI in responsible ways that protect their firm's—and their customers'—best interests.





November 2022 will go down as a watershed time in business history. Silicon Valley startup company OpenAI publicly launched its generative AI system, GPT 3.5, to power ChatGPT. Within five days, one million people downloaded the chatbot. In two months, that number grew to 100 million. The **uptake is astonishing**, even in the digital era. It took social media juggernaut TikTok nine months to reach that pinnacle and Instagram 2.5 years.ⁱ

Competitive offerings from Google (Bard), Microsoft (which licensed OpenAI's technology), Anthropic (which had **raised more than \$1.4 billion** by Augustⁱⁱ) and dozens of specialized generative AI tools have followed suit, including Amazon Web Services' Bedrock, which enables organizations to build generative AI systems through an application programming interface (API). They've unleashed **a groundswell of excitement** and rapid adoption across the world, from the paneled walls of Fortune 500 companies to the laptop computers of high school students. With billions in investment dollars from venture capital and large companies pouring in, generative AI developers are hard at work. It's a 21st century digital gold rush.

The result is a proliferation of new LLMs—and software tools to get value from them. Software companies of every type have flocked to the field, from the largest (Microsoft, Oracle, Salesforce, SAP, etc.) to startups that recently deposited a venture capital firm's check. Even still, the number of these firms is outpaced by the number of ideas about where to use generative AI.

Nonetheless, in the intensely regulated financial services and insurance world, the exuberance must be balanced with due diligence. Generative AI's promise—personal and organizational productivity improvement at scale—must be weighed against the risks: operational (privacy and security, chief among them), financial, reputational and regulatory.

Clearly, generative AI presents a major opportunity for financial institutions to increase revenue, decrease cost, shrink cycle times and reduce errors. Imagine onboarding new customers in minutes rather than days, weeks or months (in the case of corporate

customers). Think about boosting customer satisfaction because service reps can issue just-in-time, hyper-personalized answers to the daily onslaught of customer queries. Consider the impact on application development and enhancement when the technology speeds builds, minimizes defects and streamlines problem resolution.

There has never been a technology like generative AI that could write cogent prose in response to conversational human inputs. Generative AI systems could be trained to help a novice sales agent create a highly tailored pitch based on the personality, age and financial details of a prospect, and deliver the advice of a seasoned sales professional.

Generative AI can also help software developers accelerate their work. They can use it to write highly optimized code in a variety of popular programming languages. Software developers can also automate the time-intensive, painstaking collection and normalization of digital data. Through emerging APIs, they can tap structured data (e.g., from an Oracle database) and unstructured data (e.g., digitized images or posts from LinkedIn) from a variety of sources. They can then use generative AI to write code that resolves the semantic and syntactic differences that in the past would prevent an application from using the data.

Many generative AI systems will likely need to expand the number of sources from which they collect data. Similar to the way **humans learn from experiences** and expand their knowledge, the algorithms that power generative AI systems will identify new requirements and opportunities over time—but only if they are trained with data that is relevant, accurate and unbiased.

Public pronouncements of gen AI

Numerous financial institutions have publicly announced generative AI experiments that span a variety of use cases (see Figure 1). Even more are quietly putting generative AI through its paces. They see generative AI as a tonic for creating new back- and front-office operational efficiencies that help address ongoing financial pressures. According to IT sourcing authority ISG, financial services has the highest concentration **(24%) of “transformative” use cases**. Moreover, the sector accounts for 26% of mature use cases (i.e., those with a solution in process or developed that have defined quantitative return on

- investment measures).ⁱⁱⁱ
- As a starting point, generative AI enables financial institutions to dramatically extend longstanding automation efforts in ways that were not previously possible.
- **In banks:** contract writing and legal compliance
 - **In credit cards and payments:** fraud detection, cybersecurity
 - **In wealth management:** data collection and investment analysis
 - **In insurance:** claims adjudication and underwriting

Who’s using generative AI in financial services and insurance? A look at the public pronouncements			
Financial institution	Segment	Function	Experiment
Goldman Sachs & Co.	Wealth Management	Legal/IT	Document classification and categorization; software engineering.
Morgan Stanley	Wealth management	Client service/ advisory	Locating relevant information for agents to share with clients; arming advisors with the latest investment insights gleaned from a variety of company sources.
Betterment	Wealth management	Self-directed investing	Identifying daily surplus funds in customer checking accounts and automatically sweeping excess cash into a Smart Saver money market account.
JPMorgan Chase & Co.	Retail and corporate banking; wealth management	Across the institution	Augmenting and empowering employees with AI through human-centered collaborative tools and workflows; testing and enhancing a generative AI tool called IndexGPT to enhance the way investors pick, analyze and recommend financial securities such as stocks, bonds, commodities and alternatives.
USAA	Insurance, banking and investment, and retirement products and services	Across the institution	Generating new business and enhancing the member experience. USAA foresees that differentiation in generative AI will come from its adoption and integration into experiences, and from incorporating proprietary data that is well organized and integrated for model training and tuning.
Tokio Marine, North America	Global P&C insurer	IT, customer service, marketing and research	Proofs of concept in software engineering, and drafting letters, marketing content and reports on market conditions and performance.



Lemonade, Inc.	P&C insurance	Across the company	Automating and improving over 100 Identified business processes; dozens have been prototyped; expects generative AI to impact its financials in 2024.
Mastercard	Cards & payments	Cybersecurity/client services	Comparing and Identifying cybersecurity threats; product personalization.
ABN Amro Bank	Personal and business banking	Customer service	Summarizing conversations between bank staff and customers, gathering data to assist customer service reps in answering queries and avoiding repetitive questions.
Wells Fargo & Co.	Retail banking	Customer service chatbot	Extracting meaning from text inputs via an employee pilot using Google PaLM 2.
The Travelers Companies	P&C insurance	Bond & specialty, claims, professional knowledge assistant	Processing hundreds of thousands of broker submissions using proprietary large language models, which has cut processing time from hours to minutes. This improves responsiveness to customers and distribution partners, and boosts productivity. The LLMs can also ingest legal complaints filed against insureds, highlight key liability and coverage issues, and assist in routing the cases to the best-suited defense counsel. Risk-related insights can then be incorporated back into the underwriting process. Travelers is piloting a generative AI-driven knowledge assistant, trained on thousands of pages of proprietary technical source material. This provides claim professionals with easier and faster access to accurate actionable information on technical and procedural claims matters, creating more productive interactions with customers and distribution partners.
Commonwealth Bank of Australia	Global banking	Customer service; software development	Analyzing 4,500 policy documents in real time. This helps agents answer queries faster and more accurately. Software engineers use the technology to develop code.
American Express Co.	Cards & payments	Customer intelligence	Gaining deeper insights into customer patterns and behaviors.
Capital One	Diversified banking, cards & payments	Customer experience	Experimenting with LLMs; initial usage will likely be around customer experiences.

Figure 3

Source: Cognizant (from public reports)

OpenAPI's new ChatGPT Enterprise, which adds privacy protections, data analytics, higher performance and customization, could give financial institutions more confidence to pursue more scalable applications.^{iv} AWS's Bedrock offers similar capabilities by allowing organizations to build and

deploy generative AI models that adhere to their data encryption and security policies.^v

Let's look at some of the publicly disclosed experiments in two key financial sectors.

Banking, cards & payments

Back-office teams are responsible for cleaning, vouching for and preparing reports for bank regulators. Generative AI could automate these processes end to end. One large bank we know has created a generative AI “sandbox” to experiment with applications of the technology. One of these is for the very purpose of automating much of the regulatory report-writing process, which to date (done manually) can take several weeks, or even months.

JPMorgan Chase & Co., the largest bank in the US, is experimenting with generative AI across its retail, corporate banking and wealth management units. It has more than 1,000 people in data management, over 900 data scientists (AI and machine learning experts who create new models) and 600 machine learning engineers (who write the code to put models in production).^{vi} **In his 2022 letter to shareholders**, CEO Jamie Dimon wrote: “We’re imagining new ways to augment and empower employees with AI through human-centered collaborative tools and workflow, leveraging tools like large language models, including ChatGPT.” The bank has more than 300 AI use cases in production, **up 34% over the last year**.^{vii}

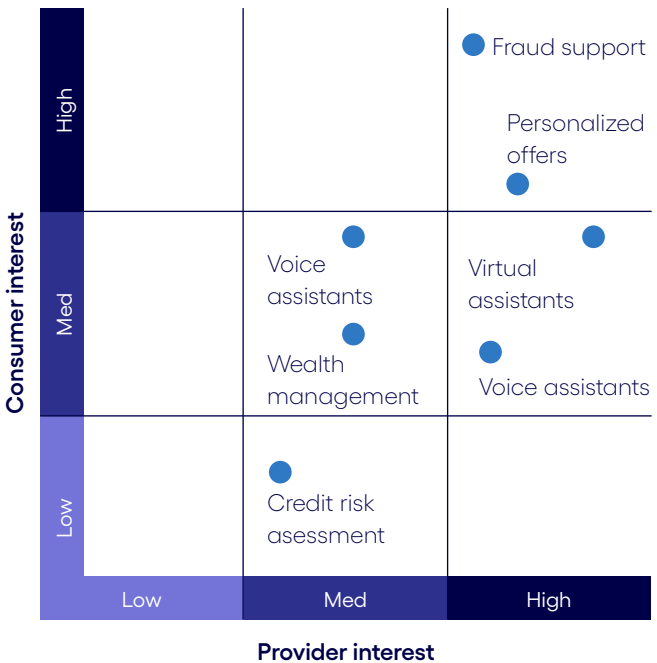
Commonwealth Bank of Australia **is using the technology** to analyze 4,500 policy documents in real time, which is helping customer service agents answer queries more quickly and accurately. Its 7,000 software engineers are also using generative AI to develop code.^{viii}

Cards and payments firms have for years used AI to detect and prevent fraud. Experimentation

with generative AI is proceeding apace. Earlier this year, Mastercard CEO Michael Miebach told investors his firm **has used generative AI** to create data sets that allow it to compare and find cybersecurity threats. It is also exploring how AI can be used in customer service.^{ix}

Analyst estimates: provider and consumer interest in generative AI applications in banking

Over a 3-year horizon



g280400

eMarketer | Insider Intelligence.com

Figure 2
Source: Insider Intelligence





Capital markets and wealth management

Capital markets and wealth management are complex businesses. They must balance multiple investment variables such as long- and short-term goals, tax implications and regulatory concerns. As a result, determining where to use generative AI is not straightforward.

This is why companies in these sectors are starting small. They are exploring whether generative AI could help investment analysts find and summarize answers to marketing and investment questions from internal and external sources. Areas in focus include earnings call transcripts, personalized client reports and routine client communications.

Goldman Sachs & Co. has **numerous generative AI proofs of concept in place**, according to Marco Argenti, chief information officer.^x One is aimed at document classification and categorization. Generative AI would take documents like legal contracts for financial products that the company receives (e.g., a bond, a loan, a derivative) and make them understandable so Goldman professionals can take informed action. Goldman has used earlier generations of AI to do this, but it is now looking at LLMs as a way to take the task to the next level. The firm is also considering using generative AI to summarize earnings calls and in-house research. In its experiments with generative AI in software engineering, it has reportedly seen up to 40% of the code written by generative AI accepted by its software developers.

Earlier in 2023, Morgan Stanley said it was **testing an OpenAI-powered chatbot** with its financial advisors.^{xi} The goal is to put most of the bank's extensive library of research and data in the hands of its advisers. Its generative AI system is scanning documents, analyst commentaries and investment research, all of which resides in multiple systems. The task normally takes the firm's investment professionals **an average of 30 minutes or more**.^{xii}

Betterment, a \$50 million New York-based digital wealth management firm, **is using generative AI** to help clients monitor monthly spending. Its algorithms scan for daily surplus funds in customer checking accounts. It then automatically sweeps excess cash into a money market account.^{xiii} The company sees this as a way to help investors make their cash reserves work better.

Where environmental, social and corporate governance (ESG) is top of mind, the on-the-fly hyper-personalization of content that generative AI can create is inspiring new ideas in serving wealth management clients. For example, customers increasingly expect transparency about their investments. For instance, they want to know if a company is aiding regional conflicts or skirting child labor laws before investing in them. Generative AI can help wealth management firms deliver these types of insights in seconds.





The collective
impact:
potentially
huge productivity
gains






With examples like these, generative AI could help financial institutions generate large-scale productivity gains: a way to simultaneously cut costs, spark innovation and streamline work. According to a recent McKinsey & Co. report, generative AI could help companies across a range of industries **generate \$2.6 trillion to \$4.4 trillion in value annually**.^{xiv} McKinsey expects roughly 75% of the value to come from customer operations, marketing and sales, software engineering and R&D.

For financial services and insurance companies, McKinsey predicts generative AI will deliver \$200 billion to \$340 billion in value annually if the technology is implemented across the business—especially to create virtual assistants, automate software development and personalize content (see Figure 3).

Using generative AI in just a few functions could drive most of the technology's impact across potential corporate use cases.

● Represent - 75% of total annual impact of generative AI



Impact as a percentage of functional spend, %

Note: Impact is averaged.

*Excluding software engineering.

Source: Comparative industry service (CIS), IHS market; Oxford economics; McKinsey corporate and business functions database; McKinsey manufacturing and supply chain 360; McKinsey sales navigator, a McKinsey database; McKinsey analysis

Figure 3

Source: McKinsey & Co.

Insurers eye improvements in pricing, claims management and service automation

Property & casualty and life insurance executives have told us they are exploring generative AI to improve operations in three primary areas:

1. Underwriting and pricing, where drafting insurance policies requires gathering data from multiple sources. Generative AI can expedite underwriting processes by swiftly analyzing risk based on vast amounts of data from internal systems and external databases. This should generate more accurate quotes and accelerate policy issuance.

2. Claims management, where staff could use LLMs to retrieve and summarize information from insurance policies, and search through emails, phone transcripts, claims forms and even customer relationship management systems to inform and speed claims processing and adjudication. These efficiencies could improve customer experiences, reduce operational costs, increase accuracy and boost team productivity.

3. Customer service, where bots or generative AI-informed service agents answer customer inquiries in near real-time by summarizing all relevant customer information to the question at hand.

Lemonade, Inc., a small (\$256 million revenue) P&C insurer based in New York, is **using generative AI systems in a variety of unnamed applications**. In its Q1 2023 shareholder letter, the firm said it expects the systems to be “somewhat impactful on our financials late this year and more significantly impactful in 2024 and beyond.”^{xv}

Tokio Marine North America has proofs of concept in several areas, including application development. The company says generative AI can translate in near-real-time a stored procedure from one programming language to another. “It could take weeks for a programmer to do the same thing,” CIO Robert Pick told a reporter. It is also **exploring use cases** that speed manual processes such as drafting customer letters, reports on market conditions and performance, and marketing content.^{xvi}

McKinsey's research predicts **big things for generative AI in insurance**. It estimates that ChatGPT and its generative AI peers could increase industry revenue by \$1.1 trillion—with about \$400 billion from pricing, underwriting and promotion technology upgrades, and \$300 billion from AI-powered customer service and personalized insurance offerings.^{xvii}

The perilous 11: Key risks may be as big as the rewards

Banks, financial services and insurance companies expect big things from generative AI, as do a growing number of experts in those industries. However, the risks—especially regulatory—may be just as big as the potential productivity and revenue benefits.

From our extensive experience in designing and developing enterprise systems in these sectors for more than 25 years, and from our deep knowledge of generative AI technology, we see 11 risks as being among the most important to manage.

1. **Misplaced trust.** Just because a generative AI application can comb through millions more documents in seconds than the average company could before the technology arrived, it doesn't mean its answers are always true. In fact, the number of so-called "fabrications"—wrong answers—is a reality in these early years of the technology. And if some of the answers aren't wrong, the algorithms that drive those answers may be biased. Further, an inability to explain results generated by black-box models will surely get a company in trouble if regulators want to know why, for example, certain customer segments are getting much higher rejection rates or pricing than others. Explainability remains a work in progress. Therefore, financial institutions must rely on humans to conduct reality checks on generative AI outputs.
2. **IP infringement.** The alleged use of copyrighted materials to train public LLMs has already **spawned a raft of lawsuits**.^{xviii} In fact, the US Federal Trade Commission is investigating OpenAI for **data leaks**.^{xix}
3. **IP loss.** Even if your company is not using someone else's intellectual property without their permission, you may still be unwittingly giving competitors your proprietary information if you use a public LLM. Let's say someone in your firm types into a generative AI chatbot the following question: "This is our underwriting model for these types of assets. How is our model different from those of other insurers?" By doing so, your underwriting model, which may have been a company secret, is now in the data repository of a public LLM. You have, in effect, given other users of that public LLM the opportunity to learn about your underwriting model if they type in a similar question. OpenAI claims its ChatGPT Enterprise will address this issue by **excluding customer prompts and data from its training models**.^{xx} Other players such as Google^{xxi} have also launched enterprise versions of their LLMs, in part to help companies **protect their IP**. The challenge will be to ensure that employees only use approved platforms.
4. **Regulatory reflux.** Regulations on data privacy, generative AI and related issues are in a **state of flux** globally. This places financial institutions that operate across borders at risk.^{xxii} At most risk, it appears, are those that are noncompliant on fiduciary, data privacy and ethical dimensions, such as erroneous reporting or delivering biased services and offers. The EU's recently approved regulatory framework to **create safer and more transparent use** of all forms of AI, including generative AI, are good first steps. The rules require generative AI systems to disclose when they use other parties' content. They also stipulate that design models prevent the illegal use of content, and that they publish summaries of copyrighted data.^{xxiii} Although a recently issued U.S. Executive Order on AI requires companies to report risks that AI-powered systems contain to the federal government, it is limited in scope. However, the Executive Order contains a set of guidelines that could eventually inform U.S. regulatory action^{xxiii}
5. **Tool/vendor roulette.** Picking the right toolset and vendor with staying power is a risky proposition given the technology's embryonic state. A generative AI platform that files bankruptcy in three years is not likely to be as easy to maintain as one whose owner has a thriving business.
6. **Unsustainable advantage.** Cloud computing vendors such as AWS have the financial resources and technology infrastructure already in place to support the immense computational requirements that banking and insurance companies will need for the compute-intensive applications of generative AI. Yet as the hyperscalers expand their offerings and make their services more affordable, that means large and small companies alike may have the same ability to use these services.



7. A commoditization of core generative AI toolsets means banks and insurers of any size can have the same computational capabilities. Early gains can evaporate quickly as fast-followers use the same tools and models as early adopters do.
8. **Audacious overreach.** Overly ambitious objectives can get CEOs, CFOs and company board members to buy into large but still speculative generative AI investments. However, if the early returns don't meet their expectations, the excitement about generative AI may turn to shrugs and skepticism. Better to be conservative about the potential benefits and use self-financing mechanisms, until they start proving themselves.
9. **Proliferating "orphan code."** Generative AI enables non-techies to become programmers without any software education. You could euphemistically call it "the democratization of software engineering." Or, pessimistically, you could call it "a recipe for orphan code," or code that's created by business managers and their non-IT staff members, and is then abandoned when they leave or lose interest. (It won't likely be in their job descriptions to maintain it.) If that code is in play, it will need to be maintained and extended by corporate IT—and connected to core operational systems.
10. **Talent droughts.** The lack of people skilled in using, developing and maintaining generative AI systems makes every experiment a gamble. There just aren't enough data scientists, machine learning operations, prompt engineers and other gurus at present for even the largest financial institutions to adequately launch and execute their generative AI experiments. As generative AI matures, talent shortfalls should recede.
11. **Security vulnerabilities.** The biggest threats right now include bad actors' ability to inject code into generative AI prompts to **extract proprietary information or cause fallacious responses**.^{xxiv} Few guardrails exist around open-source LLMs, which could open Pandora's Box to the data privacy risks listed above and inadequate cloud security mechanisms that could put sensitive data at risk.
12. **Organ rejection.** Given its embryonic state, what if employees, customers or business partners reject generative AI systems? Fears of being replaced by the technology could fuel skepticism about its value. Hollywood's actors and screenwriters are already voicing these concerns. So could

loan officers, investment bankers, wealth managers, underwriters and other financial services employees.

Dealing with these risks effectively begins in system design. Generative AI will only deliver on its lofty potential if banks and insurers embrace best practices throughout the software development lifecycle. This is why leading financial institutions are experimenting cautiously.

AI and analytics: A road well-traveled

Artificial intelligence is not new to financial institutions and insurance companies. They have embraced AI as it evolved from university labs and backroom experiments to the mainstream. Even back in the late 1980s, AI at American Express **flagged potential fraud** before its employees authorized credit card transactions.^{xxv}

Banks have also been using AI-powered natural language processing (NLP) engines and sentiment analysis tools for the last few years to improve online self-service and contact center-delivered care.

Investment banking firms have for years used in-house trading models and AI algorithms that incorporate statistical analysis and machine learning techniques to select and trade securities. Insurance companies already rely on analytics and AI in product research, underwriting and risk management, fraud detection, claims management and customer operations. In fact, **at least 80% of insurers globally** are investing in intelligent automation, according to industry researcher Celent.^{xxvi}

So what's different with generative AI? Previous generations of AI were focused on specific tasks like classification, prediction, optimization and decision making using existing data. Conversely, generative AI applications use newer forms of deep-learning algorithms to generate original content that resembles human created work. As such, they are designed to respond to prompts that leverage search vector embeddings to deliver text, images, audio and software code. And as more data is transformed into these embeddings, the likelihood of better and more precise answers improves through machine learning.

Getting prepared without getting overzealous

Generative AI has the potential to substantially improve process automation and organizational productivity. But while the potential appears to be vast, the risks may be just as large—if not larger.

It's early days, and the technology is still percolating. Financial institutions and insurers should start with low-risk, high-value use cases. Don't be motivated solely by FOMO (fear of missing out). Generative AI roadmaps should simultaneously serve two purposes:

- Drive the pace and sequence in which generative AI skills are developed.
- Tangibly demonstrate the potential business value of the technology through a series of projects that grow in scope and scale as the organization's AI capability matures.

Here's an analogy: Do you want heart surgery performed by the doctor who just graduated from medical school? Or do you want somebody who's

been doing it for a while? Going too fast without taking into account all the risk could raise the odds for significant errors. Although an extended learning curve may lengthen time to market, it is the prudent way to go. If you've not done so already, set up experiments that could make material advances in personal and organizational productivity.

Generative AI's capabilities are accelerating rapidly. The gap between vendor hype and industry adoption is tightening. Let the results from your experiments and other evidence from within and outside your industry be your North Star.

Lastly, stay current with emerging safety measures. A forum created by OpenAI, Alphabet (Google), Microsoft and Anthropic to **share research** on the safe and responsible use of LLMs is worth monitoring.^{xxvii} The **voluntary commitments** AI vendors recently made to the US government to enforce safety, security and trust is another good first step.^{xxviii}

But more industry/government cooperation is needed to create checks and balances that protect financial institution reputations and keep customers whole.



About the authors



Ed Merchant

Ed Merchant is Head of Consulting, Americas within Cognizant's Banking and Financial Services (BFS) business unit. His group is responsible for advising and assisting CxOs and other senior leaders on strategy execution for technology driven operational improvement, transformation, and innovation initiatives. He has 40-plus years of experience as an engineer and technologist focused on the implementation of mature and emerging technologies—the last 27 years exclusively in banking and financial services. He has deep expertise in helping financial institutions generate tangible business value from leading-edge technologies, including big data, advanced analytics, cloud computing, and now generative AI.

Prior to joining Cognizant, Ed was Global Solution Leader for the BFS division at another major service provider. Previously he held various regional and divisional CIO roles at a top-15 global bank, during which time he also served as Global Head of Architecture for its wholesale banking group. He has also held the position of Principal at a Big 4 consulting firm, leading a large-scale systems architecture, and engineering practice focused on trading and payments platforms.

Ed has an MS degree in mechanical engineering from Fairleigh Dickinson University and a BS in industrial education and technology from Montclair State University.



Manan Gauba

Manan heads Strategy and Partnerships for Cognizant's Banking and Financial Services (BFS) business unit. In this role, he oversees innovation and solutions development, drives Cognizant BFS's go-to-market partner strategy, and helps clients uncover new business models and streamline their processes and technologies. On partnerships, Manan works closely with leading ISVs, hyperscalers, fintechs and enterprise platform vendors that serve the banking and capital markets sectors.

Over his 20-plus year career, Manan has held senior leadership positions within strategy and operations, client management, sales and business development, marketing and HR at global 2000 firms and startups across the banking, healthcare, logistics and media industries. He began his career as a software engineer and IT services consultant before initially joining Cognizant in October 2005 as a client relationship manager. He holds a B.Tech degree in mechanical engineering from the Indian Institute of Technology (Banaras Hindu University), Varanasi and an MBA from the Indian Institute of Management, Indore.

About the authors



Craig Weber

Craig is Head of Strategy for Cognizant's Insurance business unit. In this role, he oversees the development, management and communication of practice strategy and is responsible for verticalized practice offerings. Previously, he led Cognizant's Life, Annuity and Retirement and Group Insurers practices, where he helped clients address critical digital business transformation challenges and modernize their core systems.

Prior to joining Cognizant, Craig spent 16 years at Celent, a world-renowned banking and financial services industry research house. He joined Celent in 2002 as a Senior Vice President to lead the firm's insurance industry analyst team. In 2012, Craig was named Celent CEO. A recognized industry expert and thought leader, Craig has authored dozens of reports on innovation and technology-enabled transformation in the financial services industry. He earned a BA in journalism from Saint Michaels College and holds an MA in communication from Emerson College.

To connect with our generative AI experts, please email us at GenAI_FinServices@cognizant.com.

Acknowledgments

The authors would like to thank the following for their contributions to this report: Babak Hodjat, Cognizant Chief Technology Officer, AI; Troy Danka, Cognizant Capital Markets Consulting Partner & Practice Leader; Gregory Verlinden, Cognizant Associate Vice President Analytics & AI, Benelux; and Sanghosh Bhalla, Cognizant Associate Vice President, Senior Banking Consulting Principal.



Endnotes

- i. <https://www.visualcapitalist.com/threads-100-million-users/#:~:text=Ranked%20second%2C%20Open%20AI%27s%20ChatGPT,the%20potential%20consequences%20of%20AI.>
- ii. <https://news.crunchbase.com/ai-robotics/anthropic-cash-raise-sk-teleco-goog/>
- iii. https://isg-one.com/docs/default-source/default-document-library/2023-state-of-applied-generative-ai.pdf?sfvrsn=2f43d431_4 (requires registration)
- iv. <https://openai.com/enterprise>
- v. <https://www.youtube.com/watch?v=5EDOTtYmkml>
- vi. <https://reports.jpmorganchase.com/investor-relations/2022/ar-ceo-letters.htm#specific-issues>
- vii. <https://www.businessinsider.com/jpmorgan-investor-day-slides-on-tech-spending-ficc-first-republic-2023-5#the-bank-has-over-300-ai-use-cases-in-production-5>
- viii. <https://www.afr.com/companies/financial-services/cba-goes-all-in-on-generative-ai-20230524-p5daw4>
- ix. <https://seekingalpha.com/article/4597342-mastercard-incorporated-ma-q1-2023-earnings-call-transcript>
- x. <https://www.wsj.com/articles/goldman-sachs-cio-tests-generative-ai-886b5a4b>
- xi. <https://www.businesswire.com/news/home/20230314005977/en/Morgan-Stanley-Wealth-Management-Announces-Key-Milestone-in-Innovation-Journey-with-OpenAI>
- xii. <https://fortune.com/2023/03/14/morgan-stanley-testing-openai-chatgpt-gpt4-to-help-financial-advisors/#>
- xiii. <https://tearsheet.co/artificial-intelligence/betterments-launches-tool-to-optimize-cash-savings/>
- xiv. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
- xv. https://s24.q4cdn.com/139015699/files/doc_financials/2023/q1/Shareholder-Letter-Q1-2023-FINAL-5-3-2023.pdf
- xvi. <https://www.cio.com/article/647725/it-leaders-grapple-with-shadow-ai.html>
- xvii. <https://www.mckinsey.com/industries/financial-services/our-insights/insurer-of-the-future-are-asian-insurers-keeping-up-with-ai-advances>
- xviii. <https://www.nytimes.com/2023/07/15/technology/artificial-intelligence-models-chat-data.html>
- xix. <https://www.washingtonpost.com/technology/2023/07/13/ftc-openai-chatgpt-sam-altman-lina-khan/>
- xx. <https://openai.com/enterprise>
- xxi. <https://cloud.google.com/blog/products/ai-machine-learning/enterprise-ready-generative-ai-models-go-ga-in-vertex-ai>
- xxii. <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>
- xxiii. <https://www.nytimes.com/2023/03/03/business/dealbook/lawmakers-ai-regulations.html>
- xxiv. <https://www.cobalt.io/blog/prompt-injection-attacks>
- xxv. https://bizrules.info/page/art_amexaa.htm
- xxvi. <https://www.celent.com/insights/101378827>
- xxvii. <https://www.reuters.com/technology/ai-industry-leaders-create-forum-regulate-big-machine-learning-models-2023-07-26/>
- xxviii. <https://www.whitehouse.gov/wp-content/uploads/2023/07/Ensuring-Safe-Secure-and-Trustworthy-AI.pdf>



Cognizant (Nasdaq-100: CTSI) engineers modern businesses. We help our clients modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. Together, we're improving everyday life. See how at www.cognizant.com or [@Cognizant](https://twitter.com/Cognizant).

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters

280 Bishopsgate
London
EC2M 4RB, England
Tel: +44 (0)1 020 7297 7600

India Operations Headquarters

5/535, Okkiam Thoraiyakkam,
Old Mahabalipuram Road,
Chennai, 600 096 India
Tel: 1-800-208-6999
Fax: +91 (0)1 44 4209 6060

APAC Headquarters

1 Fusionopolis Link,
Level 5 NEXUS@One-North,
North Tower Singapore 138542
Phone: +65 6812 4000

© Copyright 2023, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned here in are the property of their respective owners.