



Prompting for action

How AI agents are reshaping
the future of work

Expanded capabilities, use cases and
enterprise impact from Generative AI

About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With."

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine both challenges to AI implementation and ways to address them. The Institute collaborates with an ecosystem composed of academic research groups, startups, entrepreneurs, innovators, mature AI product leaders and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in, whether you're a board member or a C-suite leader driving strategy for your organization or a hands-on data scientist bringing an AI strategy to life, the Institute can help you learn more about how organizations across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute to access the full body of our work, subscribe to our podcasts and newsletter, and join us at our meetups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

Content

Key takeaways

- AI agents are reshaping industries by expanding the potential applications of Generative AI (GenAI) and typical language models.
- Multiagent AI systems can significantly enhance the quality of outputs and complexity of work performed by single AI agents.
- Forward-thinking businesses and governments are already implementing AI agents and multiagent AI systems across a range of use cases.
- Executive leaders should make moves now to prepare for and embrace this next era of intelligent organizational transformation.

Introduction	4
AI agents: What makes them different—and why they matter	5
Multiagent AI systems: Amplifying the potential of AI agents	7
Key benefits of AI agents and multiagent AI systems: Advantages that AI agents are unlocking for organizations today	7
Transforming strategic insights: A real-world example of a multiagent AI system	8
Achieving impact through targeted use cases: How AI agents are changing industries and enterprise domains	11
Enabling new ways of working and new horizons of innovation: Implications for strategy, risk, talent, business processes and technology	13
The road ahead: What we expect as AI agents continue to evolve	15
Charting a course into the next era of organizational transformation: Recommended actions for leaders to take now	16
Get in touch & Endnotes	17

Introduction

How can we operate faster and more efficiently?

This question has always been at the forefront of strategic agendas—but Generative AI (GenAI) is helping unlock new answers. With its ability to produce novel outputs from plain-language prompts, GenAI has enabled enterprises to significantly enhance speed and productivity across a range of business tasks. However, use cases for typical language models have only just begun to show GenAI's transformative potential. In this time of rapid AI evolution, it's time to think bigger and bolder: from streamlining routine tasks to redesigning entire workflows.

Now the question for business and government leaders is becoming:

How can we rethink our business processes with GenAI?

Large language models (LLMs) and GenAI-powered tools used by most organizations today serve as *helpful assistants*: A human worker enters a prompt, GenAI quickly produces an output. However, this interaction is largely transactional and limited in scope.

What if GenAI could be more like a *skilled collaborator* that will not only respond to requests but also plan the whole process to help solve a complex need? What if GenAI could also tap into the necessary data, digital tools and contextual knowledge to orchestrate the process end to end, autonomously?

This vision is becoming a reality with the emergence of AI agents and multiagent AI systems—a powerful advancement in what's possible through human-AI partnership. Leading companies and government agencies are already seeing the value of AI agents and putting them into practice.

In this paper, we explore what makes AI agents so groundbreaking. We then reveal how they are reshaping industries, including government and public services, by enabling new use cases, enhancing automation and accelerating the future of intelligent organizational transformation.



Adapt or fall behind

At the end of 2023, nearly **1 in 6** surveyed business leaders **said GenAI had *already* transformed their businesses.¹**

AI agents: What makes them different—and why they matter

To grasp the potential value of AI agents and their role in expanding the automation horizon, it is important to understand how they differ from the language models and GenAI applications familiar to business leaders today.

AI agents are reasoning engines that can understand context, plan workflows, connect to external tools and data, and execute actions to achieve a defined goal.

While this may sound broadly like what standalone LLMs or GenAI applications can do, there are key distinctions that make AI agents significantly more powerful. (See table, page 6.) Typical LLM-powered chatbots, for example, usually have limited ability to understand multistep prompts—much less to plan and execute whole workflows from a single prompt. In essence, they conform to the “input-output” paradigm of traditional applications and can get confused when presented with a request that must be deconstructed into multiple smaller tasks. They also struggle to reason over sequences, such as compositional tasks that require consideration of temporal and textual contexts. These limitations are even more pronounced when using small language models (SLMs), which, because they are trained on smaller volumes of data, typically sacrifice depth of knowledge and/or quality of outputs in favor of improved computational cost and speed.

As a result, early GenAI use cases have mostly been limited to standalone applications such as generating personalized ads based on a customer’s search history, reviewing contracts and legal documents to identify potential regulatory concerns, or predicting molecular behavior and drug interactions in pharmaceutical research.



AI agents excel in addressing these limitations while also leveraging capabilities of domain- and task-specific digital tools to complete more complicated tasks effectively. For example, AI agents equipped with long-term memory can remember customer and constituent interactions—including emails, chat sessions and phone calls—across digital channels, continuously learning and adjusting personalized recommendations. This contrasts with typical LLMs and SLMs, which are often limited to session-specific information. Moreover, AI agents can automate end-to-end processes, particularly those requiring sophisticated reasoning, planning and execution.

AI agents are opening new possibilities to drive enterprise productivity and program delivery through business process automation. Use cases that were once thought too complicated for GenAI can now be enabled at scale—securely and efficiently.

In other words: **AI agents don’t just interact. They more effectively reason and act on behalf of the user.**

A new paradigm for human-machine collaboration

Through their ability to reason, plan, remember and act, AI agents address key limitations of typical language models.

	 Typical language models	 AI agents
Use case scope	Automate tasks	Automate entire workflows/processes
Planning	Are not capable of planning or orchestrating workflows	Create and execute multistep plans to achieve a user's goal, adjusting actions based on real-time feedback
Memory & fine-tuning	Do not retain memory and have limited fine-tuning capabilities	Utilize short-term and long-term memory to learn from previous user interactions and provide personalized responses; Memory may be shared across multiple agents in a system
Tool integration	Are not inherently designed to integrate with external tools or systems	Augment inherent language model capabilities with APIs and tools (e.g., data extractors, image selectors, search APIs) to perform tasks
Data integration	Rely on static knowledge with fixed training cutoff dates	Adjust dynamically to new information and real-time knowledge sources
Accuracy	Typically lack self-assessment capabilities and are limited to probabilistic reasoning based on training data	Can leverage task-specific capabilities, knowledge and memory to validate and improve their own outputs and those of other agents in a system

Multiagent AI systems: Amplifying the potential of AI agents

While individual AI agents can offer valuable enhancements, the truly transformative power of AI agents comes when they work together with other agents. Such multiagent systems leverage specialized roles, enabling organizations to automate and optimize processes that individual agents might struggle to handle alone.

Multiagent AI systems employ multiple, role-specific AI agents to understand requests, plan workflows, coordinate role-specific agents, streamline actions, collaborate with humans and validate outputs.

Multiagent AI systems typically involve **standard-task agents** (e.g., user interface and data management agents) working with **specialized-skill and -tool agents** (e.g., data extractor or image interpreter agents) to achieve a goal specified by a **user**.

At the core of every AI agent is a **language model** that provides a semantic understanding of language and context—but depending on the use case, the same or different language models may be used by agents in a system. This approach can allow some agents to *share knowledge* while others *validate outputs* across the system—improving quality and consistency in the process. That potential is further enhanced by providing agents with shared **short- and long-term memory resources** that reduce the need for human prompting in the planning, validation and iteration stages of a given project or use case.

This concept extends what's possible with individual AI agents by taking a team or agency approach. By decomposing a detailed process into multiple tasks, assigning tasks to agents optimized to perform the tasks, and orchestrating agent and human collaboration at each stage of the workflow, this type of system has proven much more likely to produce higher quality, faster and more trustworthy outcomes.^{2,3}

In other words: ***Multiagent AI systems don't just reason and act on behalf of the user. They can orchestrate complex workflows in a matter of minutes.***

Key benefits of AI agents and multiagent AI systems

Capability—AI agents can automate interactions with multiple tools to perform tasks that standalone language models were not designed to achieve (e.g., browsing a website, quantitative calculations).

Productivity—Whereas standalone LLMs require constant human input and interaction to achieve desired outcomes, AI agents can plan and collaborate to execute complex workflows based on a single prompt—significantly speeding the path to delivery.

Self-learning—By tapping short- and long-term contextual memory resources that are often unavailable in a pre-trained language model, AI agents can rapidly improve their output quality over time.

Adaptability—As needs change, AI agents can reason and plan new approaches, rapidly reference new and real-time data sources, and engage with other agents to coordinate and execute outputs.

Accuracy—A key advantage of multiagent AI systems is the ability to employ “validator” agents that interact with “creator” agents to test and improve quality and reliability as part of an automated workflow.

Intelligence—When agents specializing in specific tasks work together—each applying its own memory while utilizing its own tools and reasoning capabilities—new levels of machine-powered intelligence are made possible.

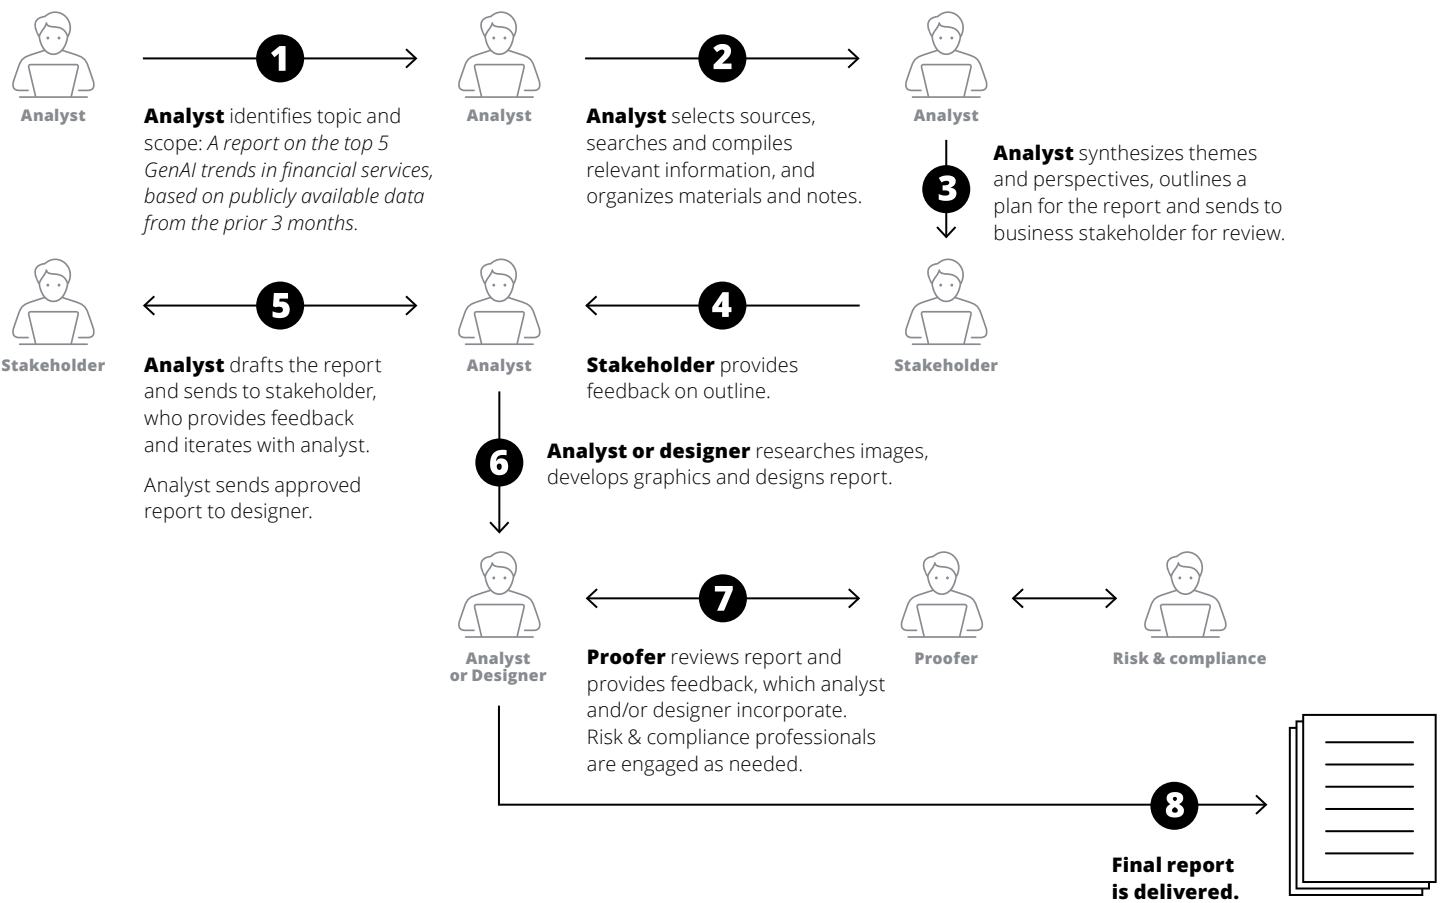
Transparency—Multiagent AI systems enhance the ability to explain AI outputs by showcasing how agents communicate and reason together, providing a clearer view of the collective decision-making and consensus-building process.

Transforming strategic insights

No matter the industry, every organization engages in research, analysis and reporting—whether about economic conditions, customer and constituent preferences, policy and pricing strategies, or other topics.

Traditionally, these projects require skilled human analysts to perform multiple steps, which can be time-consuming, utilizing research and analysis tools along with in-house subject matter expertise.

Here's what a traditional research project typically looks like.



While effective and repeatable, this approach is ...



Time-consuming

Completing a single report can take days or weeks, making it difficult to seize emerging opportunities.



Inefficient

Skilled analysts must perform many repetitive activities that take their focus away from higher-level analysis.



Difficult to scale

Companies and government agencies can struggle to hire and retain enough skilled, experienced analysts to grow their research capacity.

Deloitte has developed a multiagent AI system that can streamline and improve each step of research and reporting. **Here's how it works.**



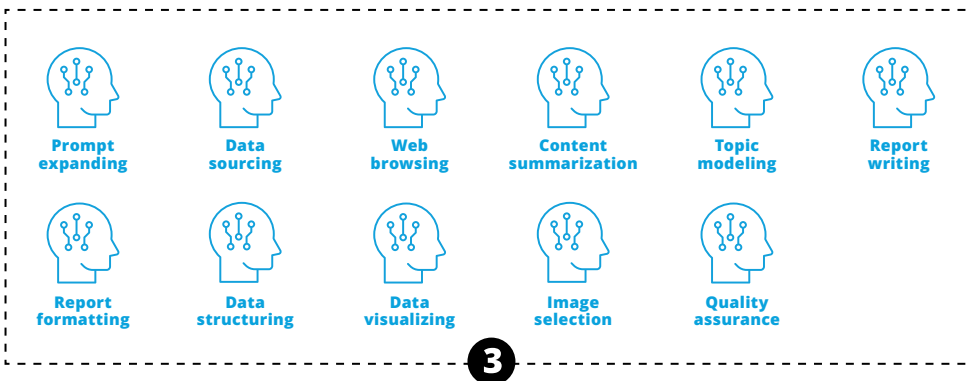

AI AGENT TYPES

Standard-task agent(s)
One or more agents that perform tasks common to all workflows

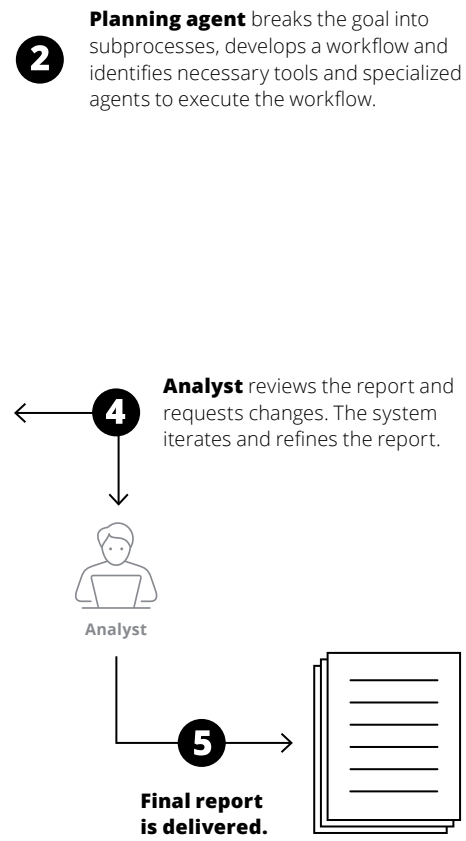
Specialized-skill & -tool agents
Role-specific agents that execute specific tasks within the workflow

All agents can access ...

- Language models (shared or separate)
- External tools & data sources as needed
- Shared short- and long-term memory



Throughout the process, the quality assurance agent checks for accuracy, quality and regulatory/brand compliance, while the data management agent ensures source materials and report iterations are documented for reference/review.



In addition to being effective and repeatable, this AI agent-powered approach is ...

- Fast**
A single, quality report can be produced in less than an hour.
- Efficient**
Skilled professionals can focus on validating, iterating and refining the report.
- Highly scalable**
In essence, this system provides an instantly available team of skilled digital workers.



Effective and efficient work depends on creativity and knowledge augmented by well-planned processes and task-appropriate tools.

That's what AI agents and multiagent AI systems can bring together.



Achieving impact through targeted use cases

Organizations across industries and sectors are already leveraging the potential of AI agents and multiagent systems to transform processes, improve efficiency, and expand impact. **Let's explore four use cases that are possible today**—two in specific industries, and two that can be applied in any business.

1 USE CASE

Individualized financial advisory and wealth management

INDUSTRY: Financial services

Financial advisory services often have relied on broad categorizations of customers based on age, income and risk tolerance. This approach can often miss the complexities of individual financial situations and goals. In today's rapidly changing financial landscape, there is an increasing demand for personalized, adaptive financial advice. Multiagent AI systems can analyze diverse data sources—including the customer's financial history, real-time market data, life events and even behavioral patterns—to help advisers create financial plans and investment strategies tailored for the specific individual. AI agents can then continuously monitor and adjust recommendations as circumstances change.

POTENTIAL ADVANTAGES ACHIEVED WITH AI AGENTS:



Hyperpersonalization

Customize financial advice to each customer's specific needs and goals, considering factors that other methods might overlook.



Continuous fine-tuning

Automatically update financial plans and strategies in response to changes in market conditions or personal circumstances.



Improved customer satisfaction

Strengthen customer relationships by providing more relevant and timely advice, leading to higher retention and satisfaction.



Enhanced scalability

Serve a larger number of customers with high-quality, personalized advice without raising costs to deliver.

2 USE CASE

Dynamic pricing and personalized promotions

INDUSTRY: Consumer

Standard pricing strategies often involve static models that do not account for real-time market conditions, customer behavior or inventory levels. Multiagent AI systems can rapidly integrate analysis based on vast amounts of real-time data—such as competitor pricing, customer purchase history and seasonal trends—to dynamically adjust prices. Additionally, they can personalize promotions based on individual customer preferences, attributes and shopping habits with the goal of improving conversion rates and elevating customer satisfaction.

POTENTIAL ADVANTAGES ACHIEVED WITH AI AGENTS:



Faster adaptation

Adjust prices instantly in response to market changes, inventory levels or customer demand—optimizing revenue.



Personalized offers

Tailor promotions to each customer's preferences and behavior, increasing the likelihood of purchase.



Greater profitability

Maximize margins and minimize discounting by optimizing pricing and promotions on an ongoing basis.

3 USE CASE

Talent acquisition and recruitment

DOMAIN: Human resources (HR)

Traditional recruitment processes often involve manual resume screening, repetitive candidate assessments and significant administrative work—which can lead to inefficiencies. AI agents can automate the end-to-end recruitment process by using natural language processing to analyze resumes, assess candidates based on skills and experience, and conduct initial screening interviews via GenAI-powered avatars. These systems can collaborate with HR professionals to ensure that qualified candidates are identified, prioritized and moved through the hiring pipeline efficiently while adhering to relevant regulations.

POTENTIAL ADVANTAGES ACHIEVED WITH AI AGENTS:



Increased efficiency

Automate tasks to allow HR teams to focus on strategic activities, shortening the time to hire.



Improved candidate matching

Analyze a broader range of data points to help match candidates to roles more accurately, improving the quality of hires.



Reduced bias

By standardizing candidate assessments and focusing on skills and experience, AI agents can help address unconscious bias in the recruitment process.



Dynamic scalability

Handle large volumes of applications, making it easier to manage hiring campaigns or recruit for multiple roles simultaneously.

4 USE CASE

Personalized customer support

DOMAIN: Customer and beneficiary service

Traditional customer and beneficiary support systems often rely on scripted interactions, which can fail to resolve complex or unique inquiries—leading to customer frustration and escalation. In contrast, multiagent AI systems can understand plain-language requests and generate relevant and natural responses that consider the customer's history, preferences and real-time context. These advanced systems can handle many complex inquiries effectively—reducing the need for escalation to live agents while improving customer/beneficiary satisfaction.

POTENTIAL ADVANTAGES ACHIEVED WITH AI AGENTS:



Greater consistency and scalability

AI agents can operate 24/7 without fatigue, maintaining a consistent quality of service no matter the volume of inquiries.



Improved customer experiences

Each customer interaction can be adjusted to individual needs, improving satisfaction and engagement.



Compounding efficiencies

The ability to learn from each interaction can help reduce response times, improve quality, and free up human service agents to focus on more nuanced customer requests.



Enabling new ways of working and new horizons of innovation

As language models continue to evolve, AI agents and systems are likely to become strategic resources and efficiency drivers for core business and government activities such as product development, regulatory compliance, customer service, constituent engagement, organizational design and others. **We see a future in which agents will transform foundational business models and entire industries, enabling new ways of working, operating and delivering value.**

That's why it's important for C-suite and public service leaders to *begin preparing now* for this next chapter in the evolution of human-machine collaboration and business innovation.

Let's explore some of the new ways of thinking and leading that should be considered during this time of rapid change.

Strategy implications

Leaders should begin integrating AI agents and multiagent AI systems into their overall strategies and future road maps. This involves reimagining business processes, investing in AI capabilities, and fostering cultures of innovation. Organizations should develop their own clear road map for AI agent adoption, identifying key areas where they can drive the most value and impact on broader business goals.

Effective change management will be crucial for successful integration. Leaders should think carefully through how they will address organizational resistance, provide training, and ensure that employees understand the value and benefits of AI agents. This includes developing a comprehensive communication strategy to keep employees and other stakeholders informed and engaged throughout the adoption process.

FOCUS AREAS

- Identify and prioritize business and service areas where AI agents can have the most immediate and measurable impact.
- Develop robust training programs to help employees understand and use AI agents in ways that improve productivity and efficiency.

Risk implications

AI agents introduce new risks that necessitate robust security and governance structures. A significant risk is potential bias in AI algorithms and training data, which can lead to inequitable decisions. Additionally, AI agents can be vulnerable to data breaches and cyberattacks, compromising sensitive information and data integrity. The complexity of AI systems also presents the risk of unintended consequences due to AI agents behaving unpredictably or making decisions not aligned with organizational goals.

To manage these risks, it is important to set clear parameters for agent interactions, monitor operational metrics, and continually ensure data ethics, privacy, security and integrity. As AI agents are integrated into core business processes, an enterprisewide governance framework with guidelines on data usage, ethics and security can further help mitigate risks. This framework should ensure compliance with relevant regulations and include continuous monitoring of AI agent interactions. Advanced security measures, such as encryption and multifactor authentication, can help protect against data breaches and cyberattacks. Training and awareness programs for employees can provide an additional defense by helping employees understand the ethical and operational considerations of working with AI agents.

FOCUS AREAS

- Identify brand and operational risks that may arise around data usage, AI agent interactions with each other and with tools, and ethics.
- Ensure model outputs are effectively tested and validated.
- Implement an AI agent governance framework that is regularly reviewed and updated as AI technologies evolve.
- Monitor emerging risks specific to AI agents such as “agent autonomy”—i.e., the risk of unintended consequences when agents make decisions with minimal human oversight.

Talent implications

The implementation of AI agents is likely to change the traditional workforce structure. As AI agents take over routine and lower-value tasks, there will likely be a high demand for human skills related to designing, implementing and operating these systems. Leaders should think through what new roles, job descriptions and job architectures are involved in building out the capability and then how to identify, recruit, train and retain this specialized talent.

Beyond the implications for tech talent, enterprise leaders should be ready to help employees across a wide variety of roles learn how to work *with* AI agents and even identify new use cases where they could improve processes. Deployed and managed well, AI agents can open up new realms of potential for human-machine collaboration—but that potential depends on workers understanding, embracing and being able to perform new roles.

FOCUS AREAS

- Communicate the benefits of AI agents, and help employees adapt to new ways of working.
- Foster a culture of innovation and continuous learning. Leaders should instill a mindset of innovation and adaptability related to AI agents.
- Explore a redesign of job architectures, workflows and performance metrics to reflect the new reality of humans and AI agents working in tandem.

Business process implications

AI agents and multiagent AI systems demand careful human evaluation of business processes—sometimes from the ground up. While agents will redefine many core processes over time, AI agents can be integrated into existing operating models today, enhancing the efficiency of current processes without the need for complete system overhauls. This approach makes it easier for organizations to adopt lower-risk agent solutions incrementally—but requires careful planning, management and alignment to ensure that AI agents are improving what people and/or other technology solutions already do well.

In use cases where AI agents do make sense to implement, human involvement will remain vitally important for tasks requiring judgment, validation and critical decision-making. This collaboration is important to help ensure that AI outputs are accurate, reliable and effective. In this paradigm, everyone working with AI agents serves as a manager—giving orders (via prompts), clarifying requests, monitoring progress, reviewing outputs and requesting or making changes as necessary.

FOCUS AREAS

- Ensure that where agents are implemented into existing business processes, those processes remain effective while driving greater efficiency and value.
- Establish processes for continuously monitoring and improving the performance of AI agents. This includes collecting and analyzing data on the performance of AI agents, identifying opportunities for improvement, and making changes as needed to optimize their performance.



Technology and data implications

Implementing AI agents can be costly, requiring substantial investment in technology and infrastructure. Organizations should carefully evaluate the value proposition and return on investment; and develop a phased approach to use cases, with a focus on “low-hanging fruit” (i.e., simpler use cases) that can lay the groundwork for more complex activations.

Quality data is the foundation for AI agents to work effectively. If data is inaccurate, incomplete or inconsistent, the agents’ outputs and actions may be unreliable or incorrect—creating both adoption and risk issues. It’s therefore essential to invest in robust data management and knowledge modeling.

Adopting trustworthy AI practices is a key to mitigating risks and ensuring ethical deployment. This includes developing AI agent solutions that are fair, transparent and accountable, and addressing potential biases in AI models.

FOCUS AREAS

- Put the right technology infrastructure in place to support the adoption and implementation of AI agents (e.g., AI orchestration platforms and scalable data lakes).
- Ensure data is properly organized, up to date and accessible to AI agents. This includes having well-defined data governance policies and procedures as well as continuous access to real-time data feeds to enable dynamic, accurate decisions.
- Establish processes for monitoring and managing the performance and ethics of AI agents and multiagent AI systems. Without transparent and trustworthy AI, customer trust and regulatory compliance are at risk.


The road ahead

The era of AI agent collaboration is still in its early stages. Interest is growing among businesses and technology providers, but comprehensive solutions are not yet common. There is much technical work to be done—particularly in terms of the reasoning and planning capabilities that will enable AI agents.

Improvements are likely to come fast. In recent months GenAI tools have shown significant improvements in reasoning and agent orchestration capabilities. Many venture capital firms are investing heavily across the spectrum of AI agent-related technologies, as are many of today’s leading GenAI and technology providers. What is available today is only a glimpse of what’s to come. Indeed, we anticipate a significant evolution of core language models, AI agents, and agent orchestration platforms within the next 12 months.

Future-focused leaders aren’t waiting on the sidelines.

Across industries, companies are already designing, testing and—in some cases—implementing agents.



Charting a course into the next era of organizational transformation

AI agents and multiagent AI systems represent more than just technological advancements. They represent a fundamental shift in how organizations can automate processes, improve human-machine collaboration, generate insights and respond dynamically to complex challenges. They offer the potential to unlock significant value across a wide range of functions—from enhancing customer interactions and optimizing supply chains to driving innovation in product development and service delivery.

The journey to realizing these benefits requires deliberate planning, strategic investments, and a commitment to fostering a culture that embraces continuous improvement and technological advancement. By aligning AI agent initiatives with core business goals, investing in the right infrastructure and nurturing a culture of innovation, your organization can be well-positioned to lead in this new era of AI-powered business transformation.

Now is the time to move.

GenAI tools are evolving rapidly—and that evolution is unlikely to slow down in the next few years. Similarly, AI agents are already being implemented by companies across industries as well as by major technology providers. So, it's important to begin exploring initial applications / use cases of agents, while setting the stage for future foundational business transformation.

To begin your own organization's journey, consider these actions:

1 Assess and prioritize use cases

Begin with a comprehensive assessment of your current operations to identify high-impact areas where AI agents can add value. Focus on processes that are ripe for automation, involve complex decision-making and/or require rapid adaptability. Prioritize these use cases to achieve quick wins and demonstrate tangible value.

2 Develop a strategic AI agent road map

Align your AI initiatives with broader business and mission objectives by creating a detailed road map that outlines the integration of AI agents into your operations. This plan should include clear milestones, timelines and success metrics to guide the deployment of AI agent-powered capabilities across the organization.

3 Invest in infrastructure and human talent development

Identify and build the necessary infrastructure to support AI agents, including scalable cloud platforms, advanced data analytics tools and robust cybersecurity measures. Simultaneously, invest in upskilling your workforce, focusing on technical skills and the ability to collaborate effectively with AI agents and multiagent systems. A well-prepared workforce is key to realizing the full transformation potential of AI agents.


4 Implement strong data governance and risk management

As AI agents become integral to your operations, it's important to establish strong governance frameworks to manage the associated risks. Implement policies that ensure data integrity, security and ethical use, while continuously monitoring AI interactions to safeguard against biases and unintended consequences. And compliance with regulatory standards should always be a top priority.

5 Nurture a culture of innovation

Experimentation and continuous learning are vital to your success. Empower your teams to explore new applications of GenAI, iterating on initial deployments to drive ongoing improvements. By embedding innovation into the fabric of your organization, you can maintain a competitive edge in a rapidly changing business environment.

Get in touch



Vivek Kulkarni
Managing Director,
AI Transformation
Deloitte LLP
vivkulkarni@deloitte.com



Scott Holcomb
Principal,
GenAI Transformation Leader
Deloitte Consulting LLP
sholcomb@deloitte.com



Prakul Sharma
Principal,
AI & Data
Deloitte Consulting LLP
praksharma@deloitte.com



Ed Van Buren
Principal,
GPS Applied AI Leader
Deloitte Consulting LLP
emvanburen@deloitte.com



Caroline Ritter
Manager,
AI Transformation
Deloitte LLP
carritter@deloitte.com

Contributors to this report:

Jim Rowan, Parth Patwari, Rajib Deb, Brijraj Limbad, Hye Ra Moon

Endnotes

1. Deborshi Dutt, Beena Ammanath, Costi Perricos and Brenna Sniderman, *Now decides next: Insights from the leading edge of generative AI adoption*, Deloitte, January 2024, p. 8, <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-state-of-gen-ai-report.pdf>, accessed September 16, 2024.
2. KaShun Shum, Shizhe Diao and Tong Zhang, *Automatic Prompt Augmentation and Selection with Chain-of-Thought from Labeled Data*, Cornell University, February 27, 2024, <https://arxiv.org/abs/2302.12822>, accessed September 16, 2024.
3. Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer and Huan Sun, *Towards Understanding Chain-of-Thought Prompting: An Empirical Study of What Matters*, Cornell University, June 1, 2023, <https://arxiv.org/pdf/2212.10001>, accessed September 16, 2024.

Prompting for action:
A series on AI agents and multiagent AI systems

[Learn key insights](#) to help guide your organization's agent-enabled journey of transformation.



Deloitte.

About Deloitte

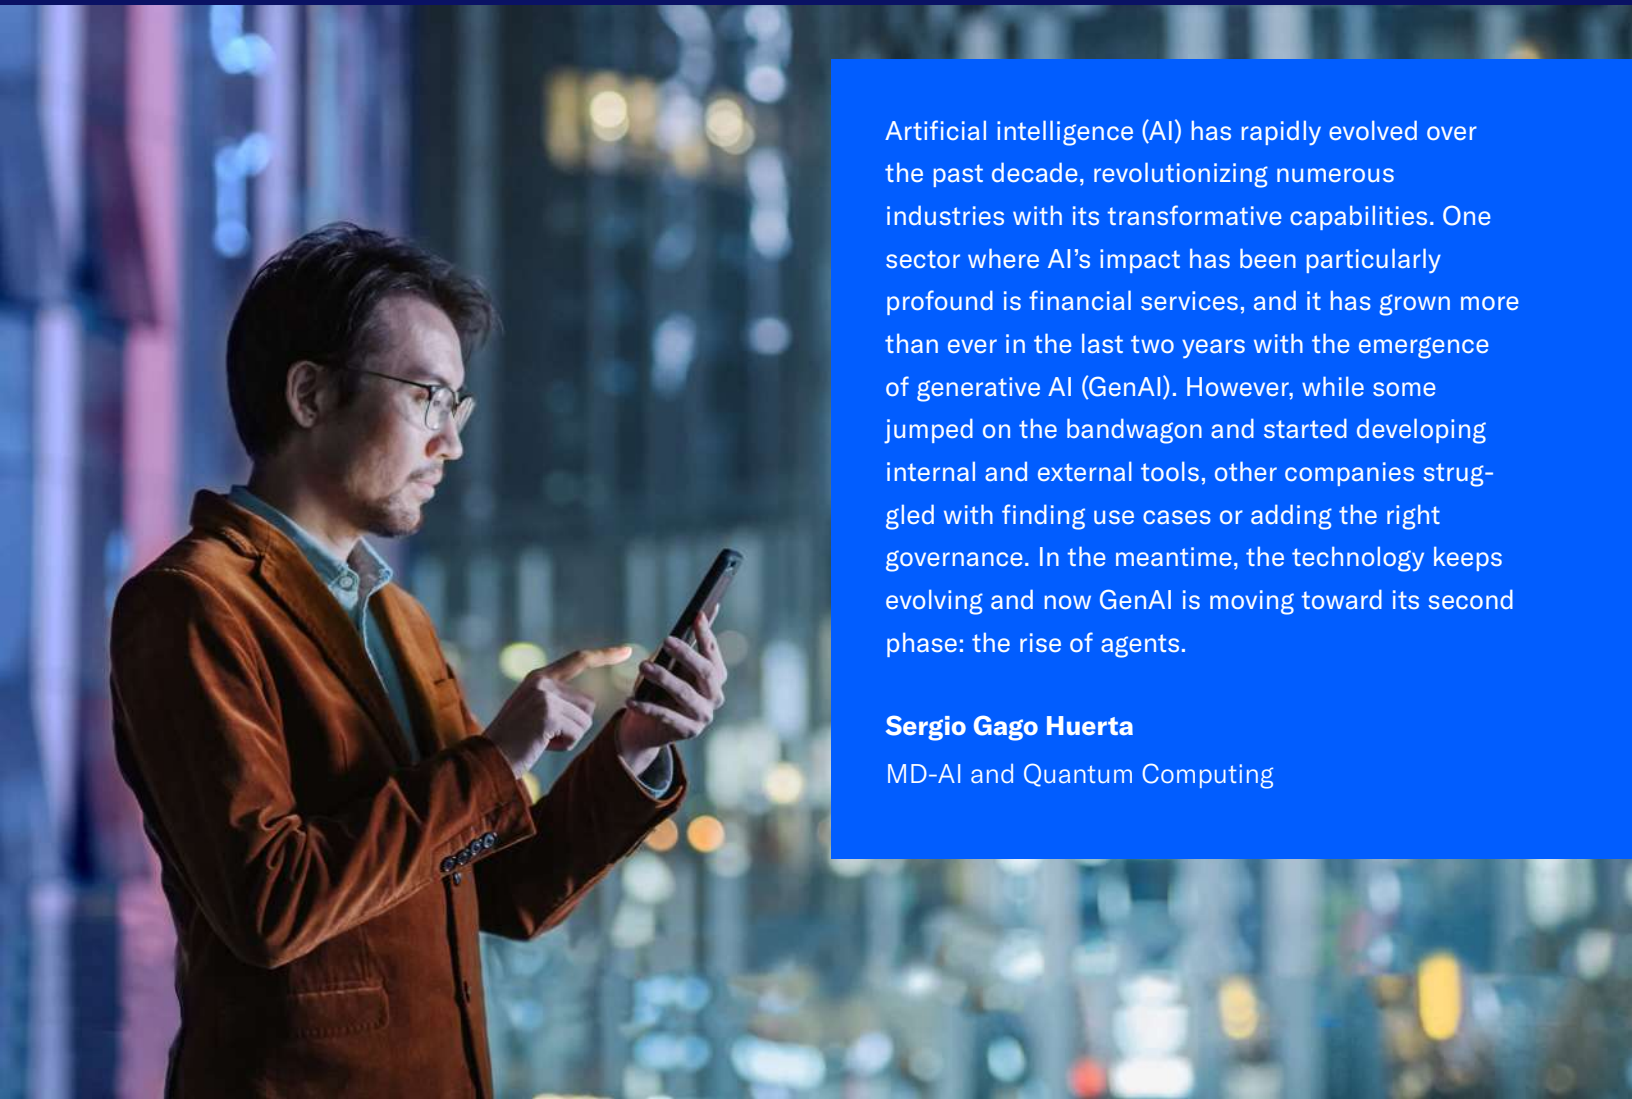
As used in this document, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Deloitte provides industry-leading audit and assurance, tax and related services, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 457,000 people worldwide make an impact that matters at www.deloitte.com.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

MOODY'S

GenAI's transformative potential in the financial sector: the evolution of agents

A man with glasses and a brown jacket is looking at a smartphone. The background is a blurred city street at night with colorful lights.

Artificial intelligence (AI) has rapidly evolved over the past decade, revolutionizing numerous industries with its transformative capabilities. One sector where AI's impact has been particularly profound is financial services, and it has grown more than ever in the last two years with the emergence of generative AI (GenAI). However, while some jumped on the bandwagon and started developing internal and external tools, other companies struggled with finding use cases or adding the right governance. In the meantime, the technology keeps evolving and now GenAI is moving toward its second phase: the rise of agents.

Sergio Gago Huerta

MD-AI and Quantum Computing

MOODY'S

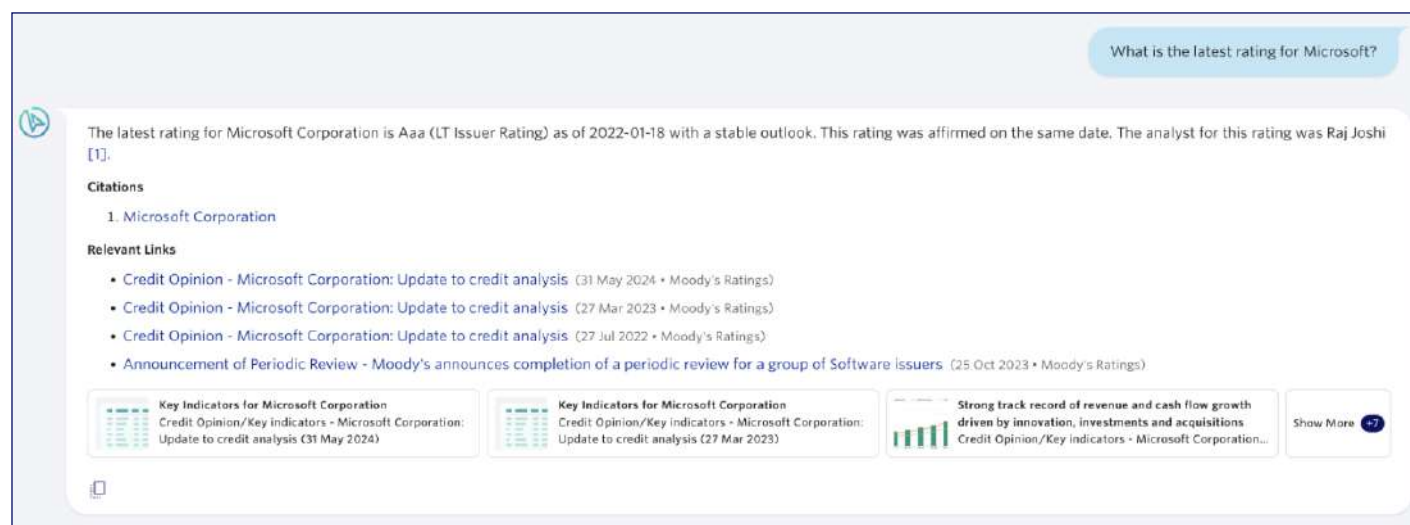
Moody's saw very early on that this technology would be a game-changer for the financial industry, and since then we have been on a GenAI journey involving creating a set of skills, assistants, and navigators and being the first financial services company to launch a GenAI tool in the market, the Moody's Research Assistant. We have also been busy researching AI agents since we believe they could potentially create a new standard for excellence in financial analytics.

But before diving into Moody's achievements, it's essential to grasp what AI agents are and how they differ from conventional chatbots.

Traditional chatbots operate based on predefined rules and scripts. They follow a set of programmed instructions to respond to user queries, often resulting in limited interactions. For instance, if a user asks a rule-based chatbot about a given company rating or financial performance, the chatbot might respond with a predetermined set of responses or direct the user to a specific webpage. These cases leverage the retrieval-augmented generation (RAG) framework that combines the power of large language models (LLMs) with individual companies' datasets.

In Moody's case, our datasets include decades of research, ratings, articles, and macroeconomic information about the world. RAG is effectively a way of telling the LLM, "We don't necessarily trust your internal knowledge and the data you have been trained on, so please use this knowledge base I am giving you access to — and by the way, please make sure you provide citations and sources to every claim you make."

Figure 1 Moody's Research Assistant provides direct access to the sources used for a given answer — for example, the last rating action for a given company as well as additional details such as financials.



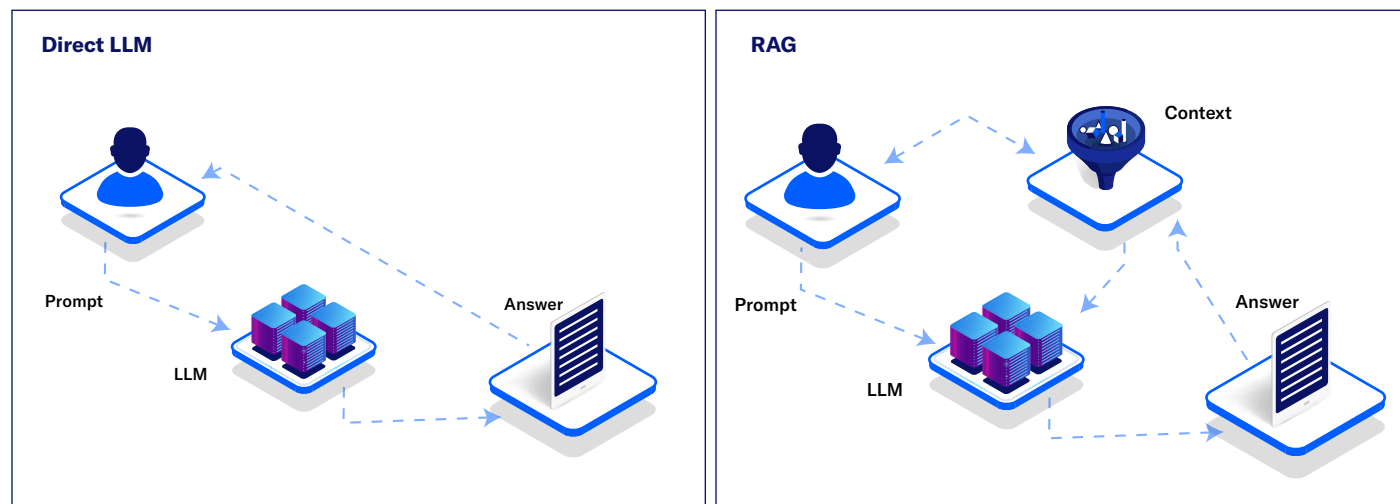
MOODY'S

We have defined four stages of GenAI maturity: assisted intelligence, augmented intelligence, automated intelligence, and autonomous intelligence.

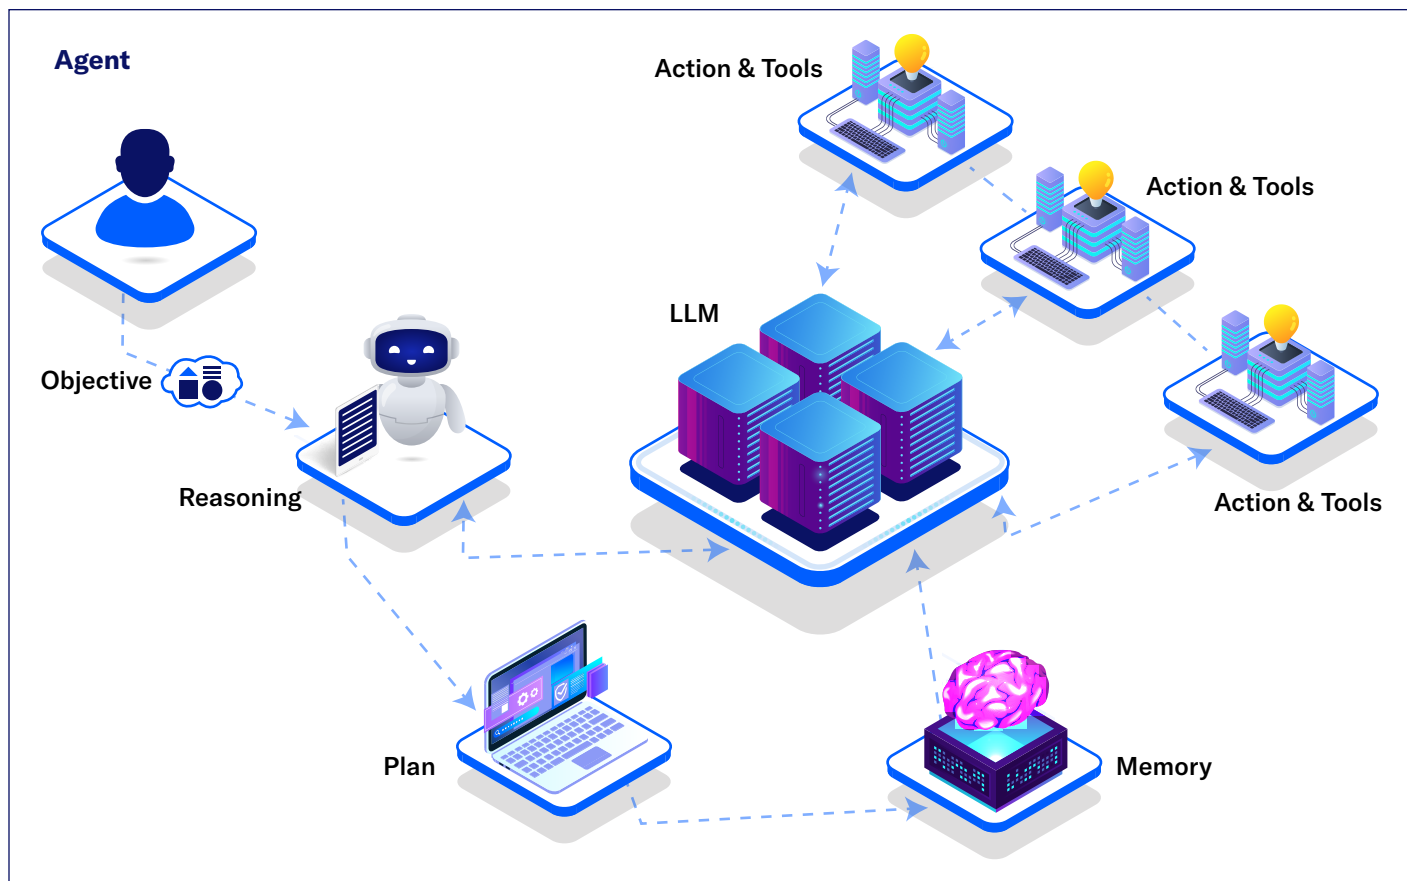
1. Initially, everyone begins with basic chatbots and rudimentary RAG frameworks, which fall under the assisted intelligence stage. At this level, AI is basic, prone to hallucinations, and requires extensive human review.
2. Augmented intelligence brings context-specific capabilities, allowing models to assess implications and determine the appropriate tools for each situation. This stage can be considered advanced RAG and tool usage, but human review and decision-making are still necessary.
3. The next step is automated intelligence, where AI can execute small tasks or make basic recommendations. Here, models are provided with extensive context, agency, and rules, along with benchmarks and quality controls. Examples of such automation include entity or address-matching, data hydration, and anomaly detection. This type of automation is similar to what classical machine learning has achieved but is now applicable to a broader domain.
4. Lastly, autonomous intelligence involves AI's ability to plan, execute those plans, evaluate outcomes, and adapt accordingly. This may involve groups of specialized agents contributing their unique perspectives and using different skills and multi-modal interaction (vision, voice, and even robotics). At this stage, we are exploring which tasks and projects are suitable. Examples include software development and analyzing a small business's financials with the same level of scrutiny applied to a Fortune 500 company.

AI agents are designed to bring us to the last stage of autonomous intelligence. They can handle more complex tasks where a simple question and answer may not be enough. Some examples could be analyzing financial reports, generating insights, and even making predictions. The agents can understand nuanced queries, learn from interactions, and provide tailored responses that evolve over time.

Figure 2 How agentic workflows really differ from RAG or pure LLM question and answer



MOODY'S



For example, imagine an AI agent designed to assist investors. Unlike a traditional chatbot that might only provide predefined sets of information, this AI agent can analyze market trends, interpret financial statements, and offer personalized investment recommendations based on the user's portfolio and risk tolerance; it can even generate a full report in the user's desired format. The agent can do that by reasoning and planning in an iterative way, whereas the chatbot would try to write the report in one go from top to bottom. Just like with human workers, the process is more efficient and yields better results when you first create a structure, research the topics, create a first draft, and have a colleague review it and provide ideas or suggestions until you have a final version. That is what agents do.

Now, imagine several LLM-powered agents that can interact with each other and have different specialties. For example, we could have a credit risk specialist agent; an environmental, social, and governance one; a news analyzer; several business analysts; and a team of coders that could create software on demand to fetch information or generate insights. These agents are the equivalent of having an army of assistants working together as a team to provide a solution for a business goal rather than just answering a question.

Moody's has been at the forefront of deploying AI agents, leveraging several advanced frameworks to enhance their functionality. During 2023 we extensively experimented with a multi-agent simulation framework. By the end of 2023 and going into 2024, we saw fundamental improvements in the technology that could lead us to build enterprise-grade production applications. On one hand, LLMs are getting better, faster, and more accurate.

MOODY'S

On the other hand, agent frameworks are now more efficient and easier to run. Below are three examples of frameworks we currently use and have based our research on:

- **Autogen:** Autogen is known for its advanced natural language processing capabilities. It allows Moody's AI agents to generate high-quality financial reports and insights by interpreting large volumes of data. For example, Autogen can automatically draft comprehensive credit reports by analyzing financial statements, market conditions, and other relevant datapoints.
- **CrewAI:** CrewAI focuses on collaborative AI applications, enhancing interactions between human analysts and AI agents. This framework is particularly useful in scenarios where human expertise and AI capabilities need to complement each other. For instance, CrewAI can assist analysts in identifying emerging market trends by analyzing data patterns and suggesting actionable insights.
- **Langgraph:** Langgraph is part of the Langchain ecosystem, employs a graph-based system to represent agent workflows, and can be used to analyze complex data relationships. Langgraph can visualize various financial indicators' interconnectedness and predict potential market shifts.

Some potential use cases where multi-agentic workflows could excel include:

Automated credit analysis: Traditionally, credit analysis involves a meticulous review of financial statements, market conditions, and qualitative factors. Powered by Autogen, Moody's AI agents streamline this process by automatically generating detailed credit reports. For instance, when assessing a company's creditworthiness, the AI agent analyzes its financial history, industry trends, and economic conditions to produce a comprehensive credit assessment. This is especially useful when you want to evaluate smaller private companies for which there is not a lot of available data.

Enhanced risk management: Managing financial risk requires constant monitoring and analysis of market conditions. Moody's uses AI agents to predict potential risks by analyzing historical data and market trends. For example, an AI agent might forecast potential economic downturns by evaluating indicators such as unemployment rates, interest rates, and geopolitical events. This requires several agents discovering the outliers and how they are correlated as well as creating different scenarios and outcomes.

Software development/service as a software: Agents can behave like authentic software development squads. One agent can be a project manager, another a quality assurance specialist, another a back-end developer, and another a front-end developer. Examples like Devin are proliferating around the world, and while it is still early days for fully fledged applications, automated agents can safely wield smaller tools. For example, queries like "Find the daily stock prices for the last two years for each ticker in my portfolio" can now be fully automated with a very easy and simple query.

Figure 3 The Autogen framework automatically created a plot for stock market price evolution.

Below we have an agent automatically plotting a trading graph with the last two years of stock prices with just one simple, unoptimized prompt — all in about 10 seconds:



Beneath the surface, the agent built software that connected to the internet, installed the right libraries, downloaded the stock prices, and printed the plot. The industry calls this “service as a software,” as opposed to software as a service (SAAS).

Figure 4 The actual source code that an agent creates and executes to produce the desired output for the user

```
primary_assistant ( to userproxy )

To plot the last 2 years of weekly stock prices for the top 5 blue-chip stocks, we will need to:

1. Identify the top 5 blue-chip stocks. This can vary depending on the criteria used to define "top" (e.g., market capitalization, performance, etc.) and the stock exchange in question. For the purpose of this task, we will assume the top 5 blue-chip stocks are based on the market capitalization from the NYSE or NASDAQ.

2. Retrieve historical stock price data for these stocks for the last 2 years on a weekly basis.

3. Plot the retrieved data.

Since I cannot browse the web or access external databases directly, I will provide you with Python code that uses the `yfinance` library to fetch historical stock data and `matplotlib` to plot the data. If you do not have these libraries installed, you will need to install them first.

Let's start by installing the required packages:
```



```
Let's start by installing the required packages:
```

```
^ hide
```

```
pip install --quiet yfinance matplotlib
```

```
After installing the packages, you can execute the following Python code to fetch and plot the stock prices:
```

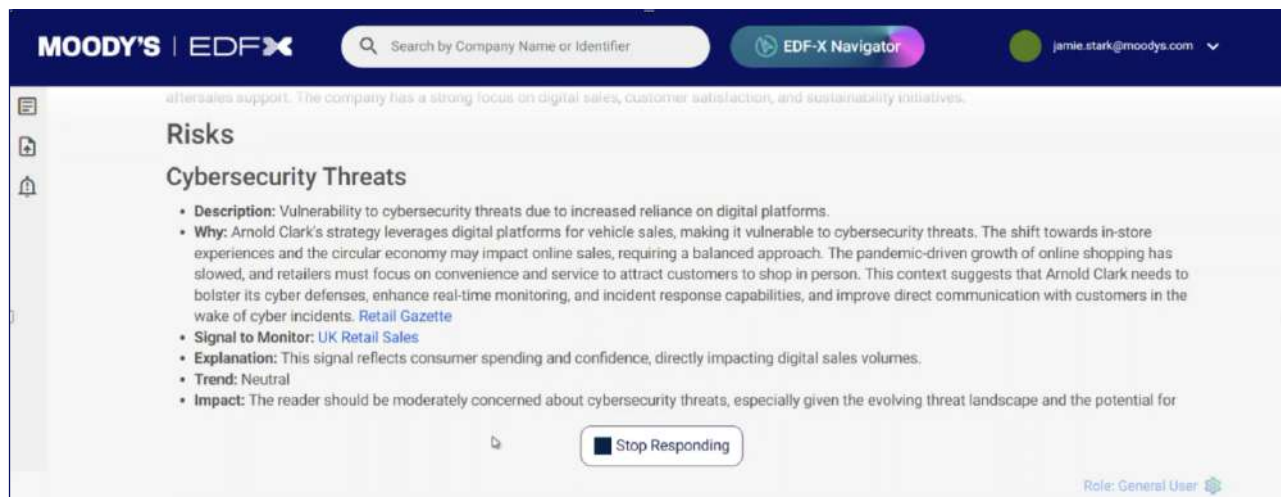
```
^ hide
```

```
# filename: plot_blue_chip_stocks.py
import yfinance as yf
import matplotlib.pyplot as plt
import pandas as pd
from datetime import datetime, timedelta
```

(These screen captures took a few seconds to produce using Autogen Studio, one of the available user interfaces for the Autogen framework.)

However, for more extensive cases in production, agents need to be tamed and their abilities constrained, from avoiding hallucinations to maintaining alignment with user requirements. Moody's has put this concept to work in one of our latest products, ReconAI, which is incorporated in our EDF-X platform. In this product, a collection of agents evaluates all the potential metrics that could trigger Early Warning Signals for any company of any size in a dynamic and customized way. For example, since we can trigger infinite agents, we can evaluate those Early Warning Signals from any public or private company and not just the Fortune 500 organizations that typically get analyzed. Another collection of agents can monitor everything that happens in the world, from news events to controversies to macroeconomic indicators to regulatory filings, and extract the signal on which companies will be affected — all in real time.

Figure 5 A screenshot of a ReconAI agent evaluating which signals to monitor for specific cybersecurity threats to a given company



MOODY'S

As a company integrates AI agents into its operations, evaluating their performance and effectiveness is crucial. There are several methods companies can take to help improve agentic-driven solutions' accuracy and reliability.

- **Retrieval-augmented generation evaluation:** As previously mentioned, RAG is a framework that combines information retrieval with generative models. Moody's uses Ragas, which is an automated evaluation of RAG with LLM as a judge — that is, another LLM is "asked" for the correctness of the answer following an extremely specific scorecard and criteria) among other frameworks to evaluate our AI agents' performance. This approach involves using LLM judges to assess the quality of generated content when reference data is unavailable. BERTScore, which measures the contextual similarity between two passages, is also applied when reference data is present to provide additional metrics on responses' relevance and accuracy.
- **Search engine/retrieval performance:** AI agents are also evaluated based on their performance in search-related tasks. Moody's leverages LLM judges to assess search results' relevance and AI agents' effectiveness in retrieving relevant information.
- **Multiple agent voting/mixture of experts:** The last method we use and experiment with is a mixture of agents and voting. Foundational models may have biases, system prompts, and guardrails that influence their answers. In some cases, they may be less confrontational. In several of our use cases, we execute the same task through several models (for example, GPT-4, Claude 2, Llama 3, Gemini, and a Moody's fine-tuned model) and hold a majority voting decision. Each answer provided by each of the underlying models gets scored and evaluated; this way we increase the accuracy and reduce the hallucination possibilities.

This is just the beginning, though. Thanks to having worked on classical machine learning, we had many of the basics covered. But this industry and technology move very fast, especially in benchmarking, which includes some of our following research lines:

- **Specialized models:** Research is underway to develop specialized models for LLM-as-a-judge, such as Google's Cappy and KAIST's Prometheus 2. These models aim to improve AI judgments' accuracy and consistency.
- **Financial benchmark datasets:** Moody's is working on creating benchmark datasets to measure LLM performance across different applications. These datasets will help detect performance drift and maintain LLM API providers' quality of service.
- **Comparative studies:** Future research will focus on comparing specialized versus general-purpose LLMs for evaluation purposes. Developing consistent scoring mechanisms for LLM judges is also a priority to enhance evaluation reliability.

As Moody's continues to lead the way in AI agent technology, several trends and developments are shaping the future of financial services, including increasing automation and depth (i.e., applying the same analysis to every company regardless of its size), enhanced personalization and contextualization, and greater integration across financial systems.

Critical technical components are in place to be leveraged by RAG and Tuning approaches

01

Data interoperability

What do you have on Tesla in all your databases?

02

Entity resolution

That's not what I meant. MCO is not the Orlando Airport.

03

Data lineage and trust

Where did this data come from? Can I trust it?

04

API architecture

Can I easily plug Moody's data into my existing workflow tools?

05

Permissions and entitlements

Who is authorized to access this data?

AI safety considerations: As we advance through the stages of GenAI maturity, it is crucial to comprehensively address AI safety. It is paramount that AI systems operate reliably and ethically.

This involves implementing robust safety protocols, continuous monitoring, and rigorous testing to prevent unintended consequences. By prioritizing AI safety, we can build trust and confidence in AI technologies and make sure they are used responsibly and beneficially.



Moody's journey in adopting and deploying AI agents exemplifies GenAI's transformative potential in the financial sector. By leveraging advanced frameworks, addressing evaluation challenges, and exploring future research directions, Moody's continues to set new benchmarks in financial analytics and risk management. As AI technology evolves, AI agents' role in enhancing financial services will become increasingly pivotal, driving innovation and delivering greater value to stakeholders.

Even though the technology had been around for several years, 2022 saw the birth of GenAI, while 2023 was the year of RAG. In 2024, we're seeing GenAI 2.0's evolution with agents. While it is still in its infancy and we have a lot of ground left to cover, the development speed is rapid and the impact is potentially exponential. Being at the forefront of applied research yields great benefits on how our industry is evolving. Moody's was the first financial institution with a GenAI product in the market with paying customers, and this year we will be the first with an agentic platform under our EDF-X suite.

By taking these steps, enterprises can deepen their understanding of AI and truly embrace the opportunities that unfold as the technology continues to evolve.



© 2024 Moody's Corporation, Moody's Investors Service, Inc., Moody's Analytics, Inc. and/or their licensors and affiliates (collectively, "MOODY'S"). All rights reserved.

CREDIT RATINGS ISSUED BY MOODY'S CREDIT RATINGS AFFILIATES ARE THEIR CURRENT OPINIONS OF THE RELATIVE FUTURE CREDIT RISK OF ENTITIES, CREDIT COMMITMENTS, OR DEBT OR DEBT-LIKE SECURITIES, AND MATERIALS, PRODUCTS, SERVICES AND INFORMATION PUBLISHED OR OTHERWISE MADE AVAILABLE BY MOODY'S (COLLECTIVELY, "MATERIALS") MAY INCLUDE SUCH CURRENT OPINIONS. MOODY'S DEFINES CREDIT RISK AS THE RISK THAT AN ENTITY MAY NOT MEET ITS CONTRACTUAL FINANCIAL OBLIGATIONS AS THEY COME DUE AND ANY ESTIMATED FINANCIAL LOSS IN THE EVENT OF DEFAULT OR IMPAIRMENT. SEE APPLICABLE MOODY'S RATING SYMBOLS AND DEFINITIONS PUBLICATION FOR INFORMATION ON THE TYPES OF CONTRACTUAL FINANCIAL OBLIGATIONS ADDRESSED BY MOODY'S CREDIT RATINGS. CREDIT RATINGS DO NOT ADDRESS ANY OTHER RISK, INCLUDING BUT NOT LIMITED TO: LIQUIDITY RISK, MARKET VALUE RISK, OR PRICE VOLATILITY. CREDIT RATINGS, NON-CREDIT ASSESSMENTS ("ASSESSMENTS"), AND OTHER OPINIONS INCLUDED IN MOODY'S MATERIALS ARE NOT STATEMENTS OF CURRENT OR HISTORICAL FACT. MOODY'S MATERIALS MAY ALSO INCLUDE QUANTITATIVE MODEL-BASED ESTIMATES OF CREDIT RISK AND RELATED OPINIONS OR COMMENTARY PUBLISHED BY MOODY'S ANALYTICS, INC. AND/OR ITS AFFILIATES. MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS AND MATERIALS DO NOT CONSTITUTE OR PROVIDE INVESTMENT OR FINANCIAL ADVICE, AND MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS AND MATERIALS ARE NOT AND DO NOT PROVIDE RECOMMENDATIONS TO PURCHASE, SELL, OR HOLD PARTICULAR SECURITIES. MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS AND MATERIALS DO NOT COMMENT ON THE SUITABILITY OF AN INVESTMENT FOR ANY PARTICULAR INVESTOR. MOODY'S ISSUES ITS CREDIT RATINGS, ASSESSMENTS AND OTHER OPINIONS AND PUBLISHES OR OTHERWISE MAKES AVAILABLE ITS MATERIALS WITH THE EXPECTATION AND UNDERSTANDING THAT EACH INVESTOR WILL, WITH DUE CARE, MAKE ITS OWN STUDY AND EVALUATION OF EACH SECURITY THAT IS UNDER CONSIDERATION FOR PURCHASE, HOLDING, OR SALE.

MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS, AND MATERIALS ARE NOT INTENDED FOR USE BY RETAIL INVESTORS AND IT WOULD BE RECKLESS AND INAPPROPRIATE FOR RETAIL INVESTORS TO USE MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS OR MATERIALS WHEN MAKING AN INVESTMENT DECISION. IF IN DOUBT YOU SHOULD CONTACT YOUR FINANCIAL OR OTHER PROFESSIONAL ADVISER.

ALL INFORMATION CONTAINED HEREIN IS PROTECTED BY LAW, INCLUDING BUT NOT LIMITED TO, COPYRIGHT LAW, AND NONE OF SUCH INFORMATION MAY BE COPIED OR OTHERWISE REPRODUCED, REPACKAGED, FURTHER TRANSMITTED, TRANSFERRED, DISSEMINATED, REDISTRIBUTED OR RESOLD, OR STORED FOR SUBSEQUENT USE FOR ANY SUCH PURPOSE, IN WHOLE OR IN PART, IN ANY FORM OR MANNER OR BY ANY MEANS WHATSOEVER, BY ANY PERSON WITHOUT MOODY'S PRIOR WRITTEN CONSENT. FOR CLARITY, NO INFORMATION CONTAINED HEREIN MAY BE USED TO DEVELOP, IMPROVE, TRAIN OR RETRAIN ANY SOFTWARE PROGRAM OR DATABASE, INCLUDING, BUT NOT LIMITED TO, FOR ANY ARTIFICIAL INTELLIGENCE, MACHINE LEARNING OR NATURAL LANGUAGE PROCESSING SOFTWARE, ALGORITHM, METHODOLOGY AND/OR MODEL.

MOODY'S CREDIT RATINGS, ASSESSMENTS, OTHER OPINIONS AND MATERIALS ARE NOT INTENDED FOR USE BY ANY PERSON AS A BENCHMARK AS THAT TERM IS DEFINED FOR REGULATORY PURPOSES AND MUST NOT BE USED IN ANY WAY THAT COULD RESULT IN THEM BEING CONSIDERED A BENCHMARK.

All information contained herein is obtained by MOODY'S from sources believed by it to be accurate and reliable. Because of the possibility of human or mechanical error as well as other factors, however, all information contained herein is provided "AS IS" without warranty of any kind. MOODY'S adopts all necessary measures so that the information it uses in assigning a credit rating is of sufficient quality and from sources MOODY'S considers to be reliable including, when appropriate, independent third-party sources. However, MOODY'S is not an auditor and cannot in every instance independently verify or validate information received in the credit rating process or in preparing its Materials.

To the extent permitted by law, MOODY'S and its directors, officers, employees, agents, representatives, licensors and suppliers disclaim liability to any person or entity for any indirect, special, consequential, or incidental losses or damages whatsoever arising from or in connection with the information contained herein or the use of or inability to use any such information, even if MOODY'S or any of its directors, officers, employees, agents, representatives, licensors or suppliers is advised in advance of the possibility of such losses or damages, including but not limited to: (a) any loss of present or prospective profits or (b) any loss or damage arising where the relevant financial instrument is not the subject of a particular credit rating assigned by MOODY'S.

To the extent permitted by law, MOODY'S and its directors, officers, employees, agents, representatives, licensors and suppliers disclaim liability for any direct or compensatory losses or damages caused to any person or entity, including but not limited to by any negligence (but excluding fraud, willful misconduct or any other type of liability that, for the avoidance of doubt, by law cannot be excluded) on the part of, or any contingency within or beyond the control of, MOODY'S or any of its directors, officers, employees, agents, representatives, licensors or suppliers, arising from or in connection with the information contained herein or the use of or inability to use any such information.

NO WARRANTY, EXPRESS OR IMPLIED, AS TO THE ACCURACY, TIMELINESS, COMPLETENESS, MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OF ANY CREDIT RATING, ASSESSMENT, OTHER OPINION OR INFORMATION IS GIVEN OR MADE BY MOODY'S IN ANY FORM OR MANNER WHATSOEVER.

Moody's Investors Service, Inc., a wholly-owned credit rating agency subsidiary of Moody's Corporation ("MCO"), hereby discloses that most issuers of debt securities (including corporate and municipal bonds, debentures, notes and commercial paper) and preferred stock rated by Moody's Investors Service, Inc. have, prior to assignment of any credit rating, agreed to pay to Moody's Investors Service, Inc. for credit ratings opinions and services rendered by it. MCO and Moody's Investors Service also maintain policies and procedures to address the independence of Moody's Investors Service credit ratings and credit rating processes. Information regarding certain affiliations that may exist between directors of MCO and rated entities, and between entities who hold credit ratings from Moody's Investors Service, Inc. and have also publicly reported to the SEC an ownership interest in MCO of more than 5%, is posted annually at www.moody.com under the heading "Investor Relations — Corporate Governance — Charter Documents - Director and Shareholder Affiliation Policy."

Moody's SF Japan K.K., Moody's Local AR Agente de Calificación de Riesgo S.A., Moody's Local BR Agência de Classificação de Risco LTDA, Moody's Local MX S.A. de C.V. I.C.V., Moody's Local PE Clasificadora de Riesgo S.A., and Moody's Local PA Clasificadora de Riesgo S.A. (collectively, the "Moody's Non-NRSRO CRAs") are all indirectly wholly-owned credit rating agency subsidiaries of MCO. None of the Moody's Non-NRSRO CRAs is a Nationally Recognized Statistical Rating Organization.

Additional terms for Australia only: Any publication into Australia of this document is pursuant to the Australian Financial Services License of MOODY'S affiliate, Moody's Investors Service Pty Limited ABN 61 003 399 657AFSL 336969 and/or Moody's Analytics Australia Pty Ltd ABN 94 105 136 972 AFSL 383569 (as applicable). This document is intended to be provided only to "wholesale clients" within the meaning of section 761G of the Corporations Act 2001. By continuing to access this document from within Australia, you represent to MOODY'S that you are, or are accessing the document as a representative of, a "wholesale client" and that neither you nor the entity you represent will directly or indirectly disseminate this document or its contents to "retail clients" within the meaning of section 761G of the Corporations Act 2001. MOODY'S credit rating is an opinion as to the creditworthiness of a debt obligation of the issuer, not on the equity securities of the issuer or any form of security that is available to retail investors.

Additional terms for India only: Moody's credit ratings, Assessments, other opinions and Materials are not intended to be and shall not be relied upon or used by any users located in India in relation to securities listed or proposed to be listed on Indian stock exchanges.

Additional terms with respect to Second Party Opinions (as defined in Moody's Investors Service Rating Symbols and Definitions): Please note that a Second Party Opinion ("SPO") is not a "credit rating". The issuance of SPOs is not a regulated activity in many jurisdictions, including Singapore. JAPAN: In Japan, development and provision of SPOs fall under the category of "Ancillary Businesses", not "Credit Rating Business", and are not subject to the regulations applicable to "Credit Rating Business" under the Financial Instruments and Exchange Act of Japan and its relevant regulation. PRC: Any SPO: (1) does not constitute a PRC Green Bond Assessment as defined under any relevant PRC laws or regulations; (2) cannot be included in any registration statement, offering circular, prospectus or any other documents submitted to the PRC regulatory authorities or otherwise used to satisfy any PRC regulatory disclosure requirement; and (3) cannot be used within the PRC for any regulatory purpose or for any other purpose which is not permitted under relevant PRC laws or regulations. For the purposes of this disclaimer, "PRC" refers to the mainland of the People's Republic of China, excluding Hong Kong, Macau and Taiwan.