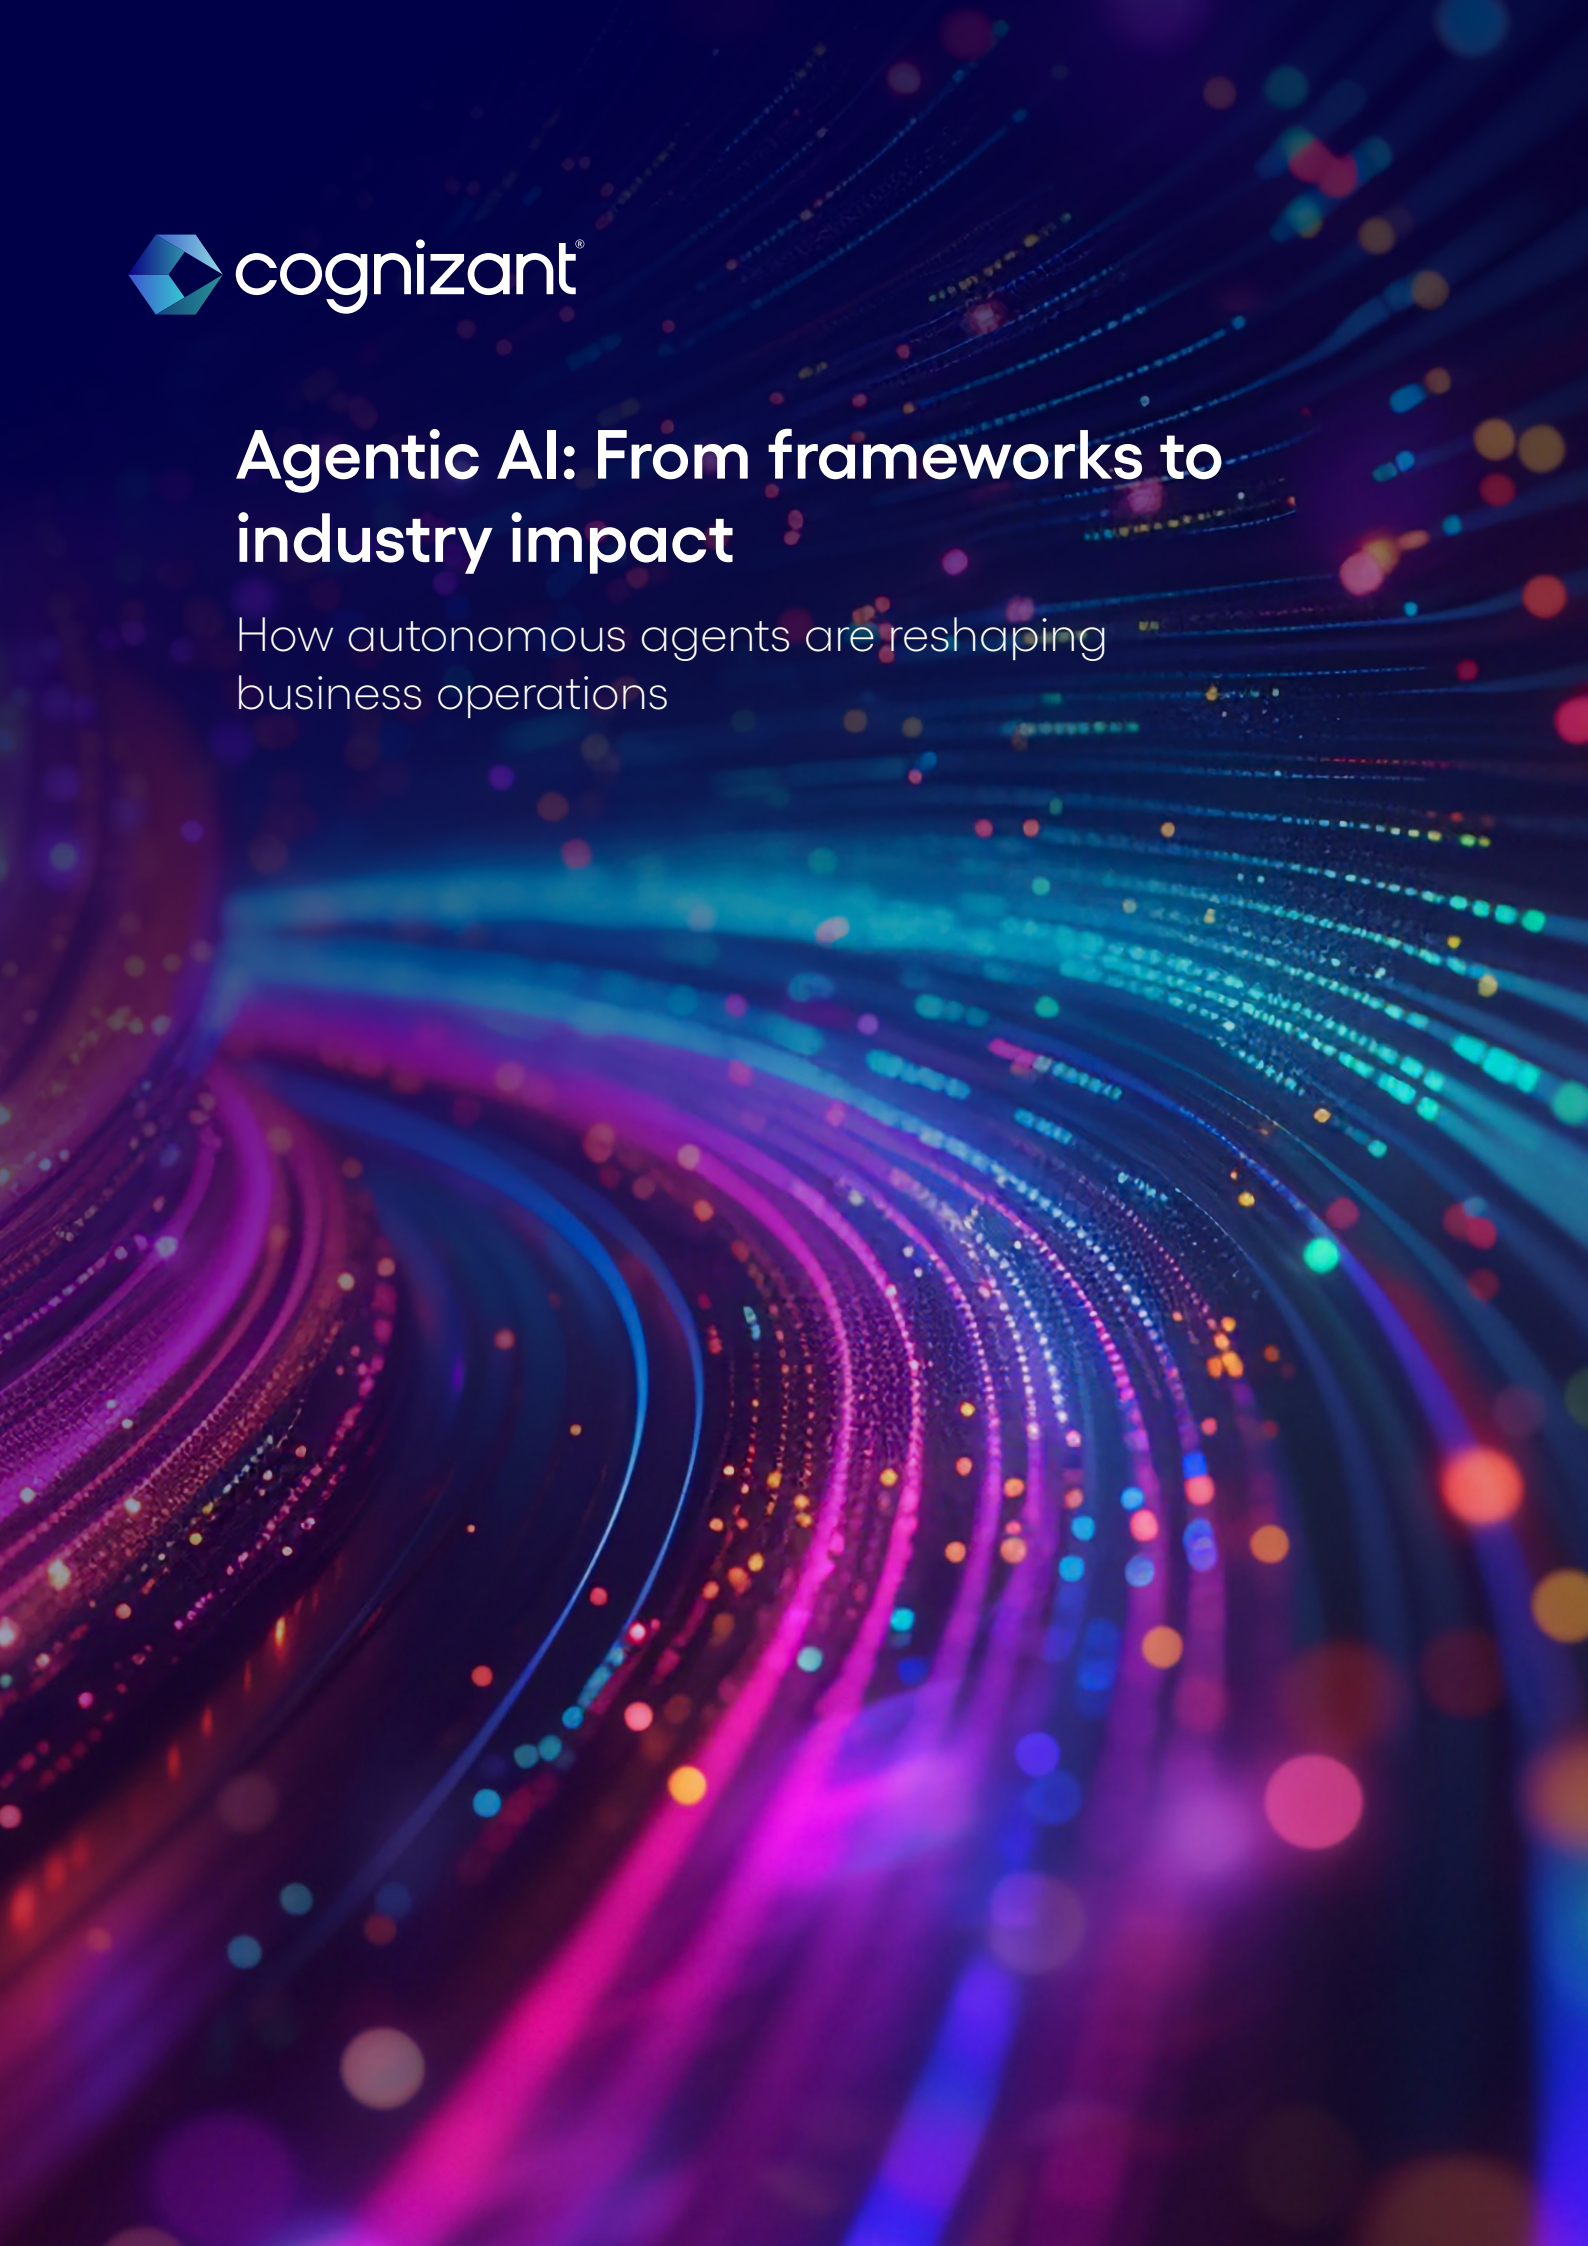
cognizant
Agentic AI: From frameworks to industry impact
How autonomous agents are reshaping business operations

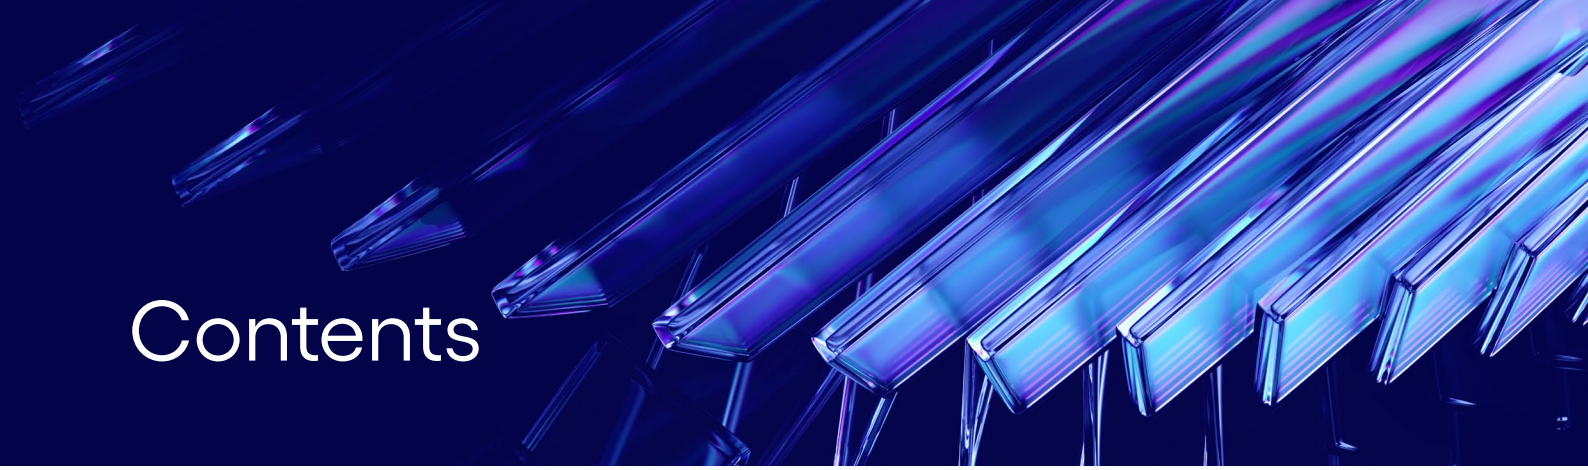

Contents

Introduction

Agentic AI refers to a class of agent-based artificial intelligence systems in which autonomous agents operate independently to achieve specific objectives. These systems are engineered to emulate human-like intelligence and behavior, enabling machines to function with a level of autonomy that was previously unattainable. This paper explores the diverse applications of agentic AI and underscores its transformative potential across industries and everyday life.

The agentic AI workflow

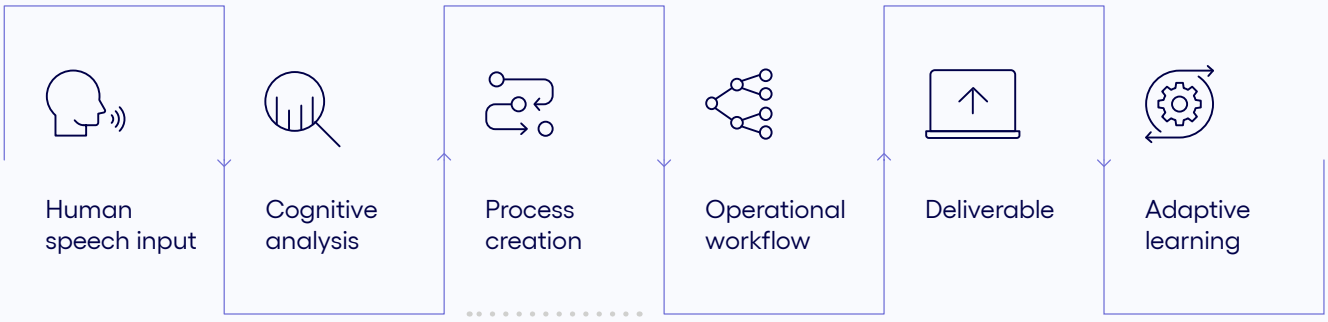

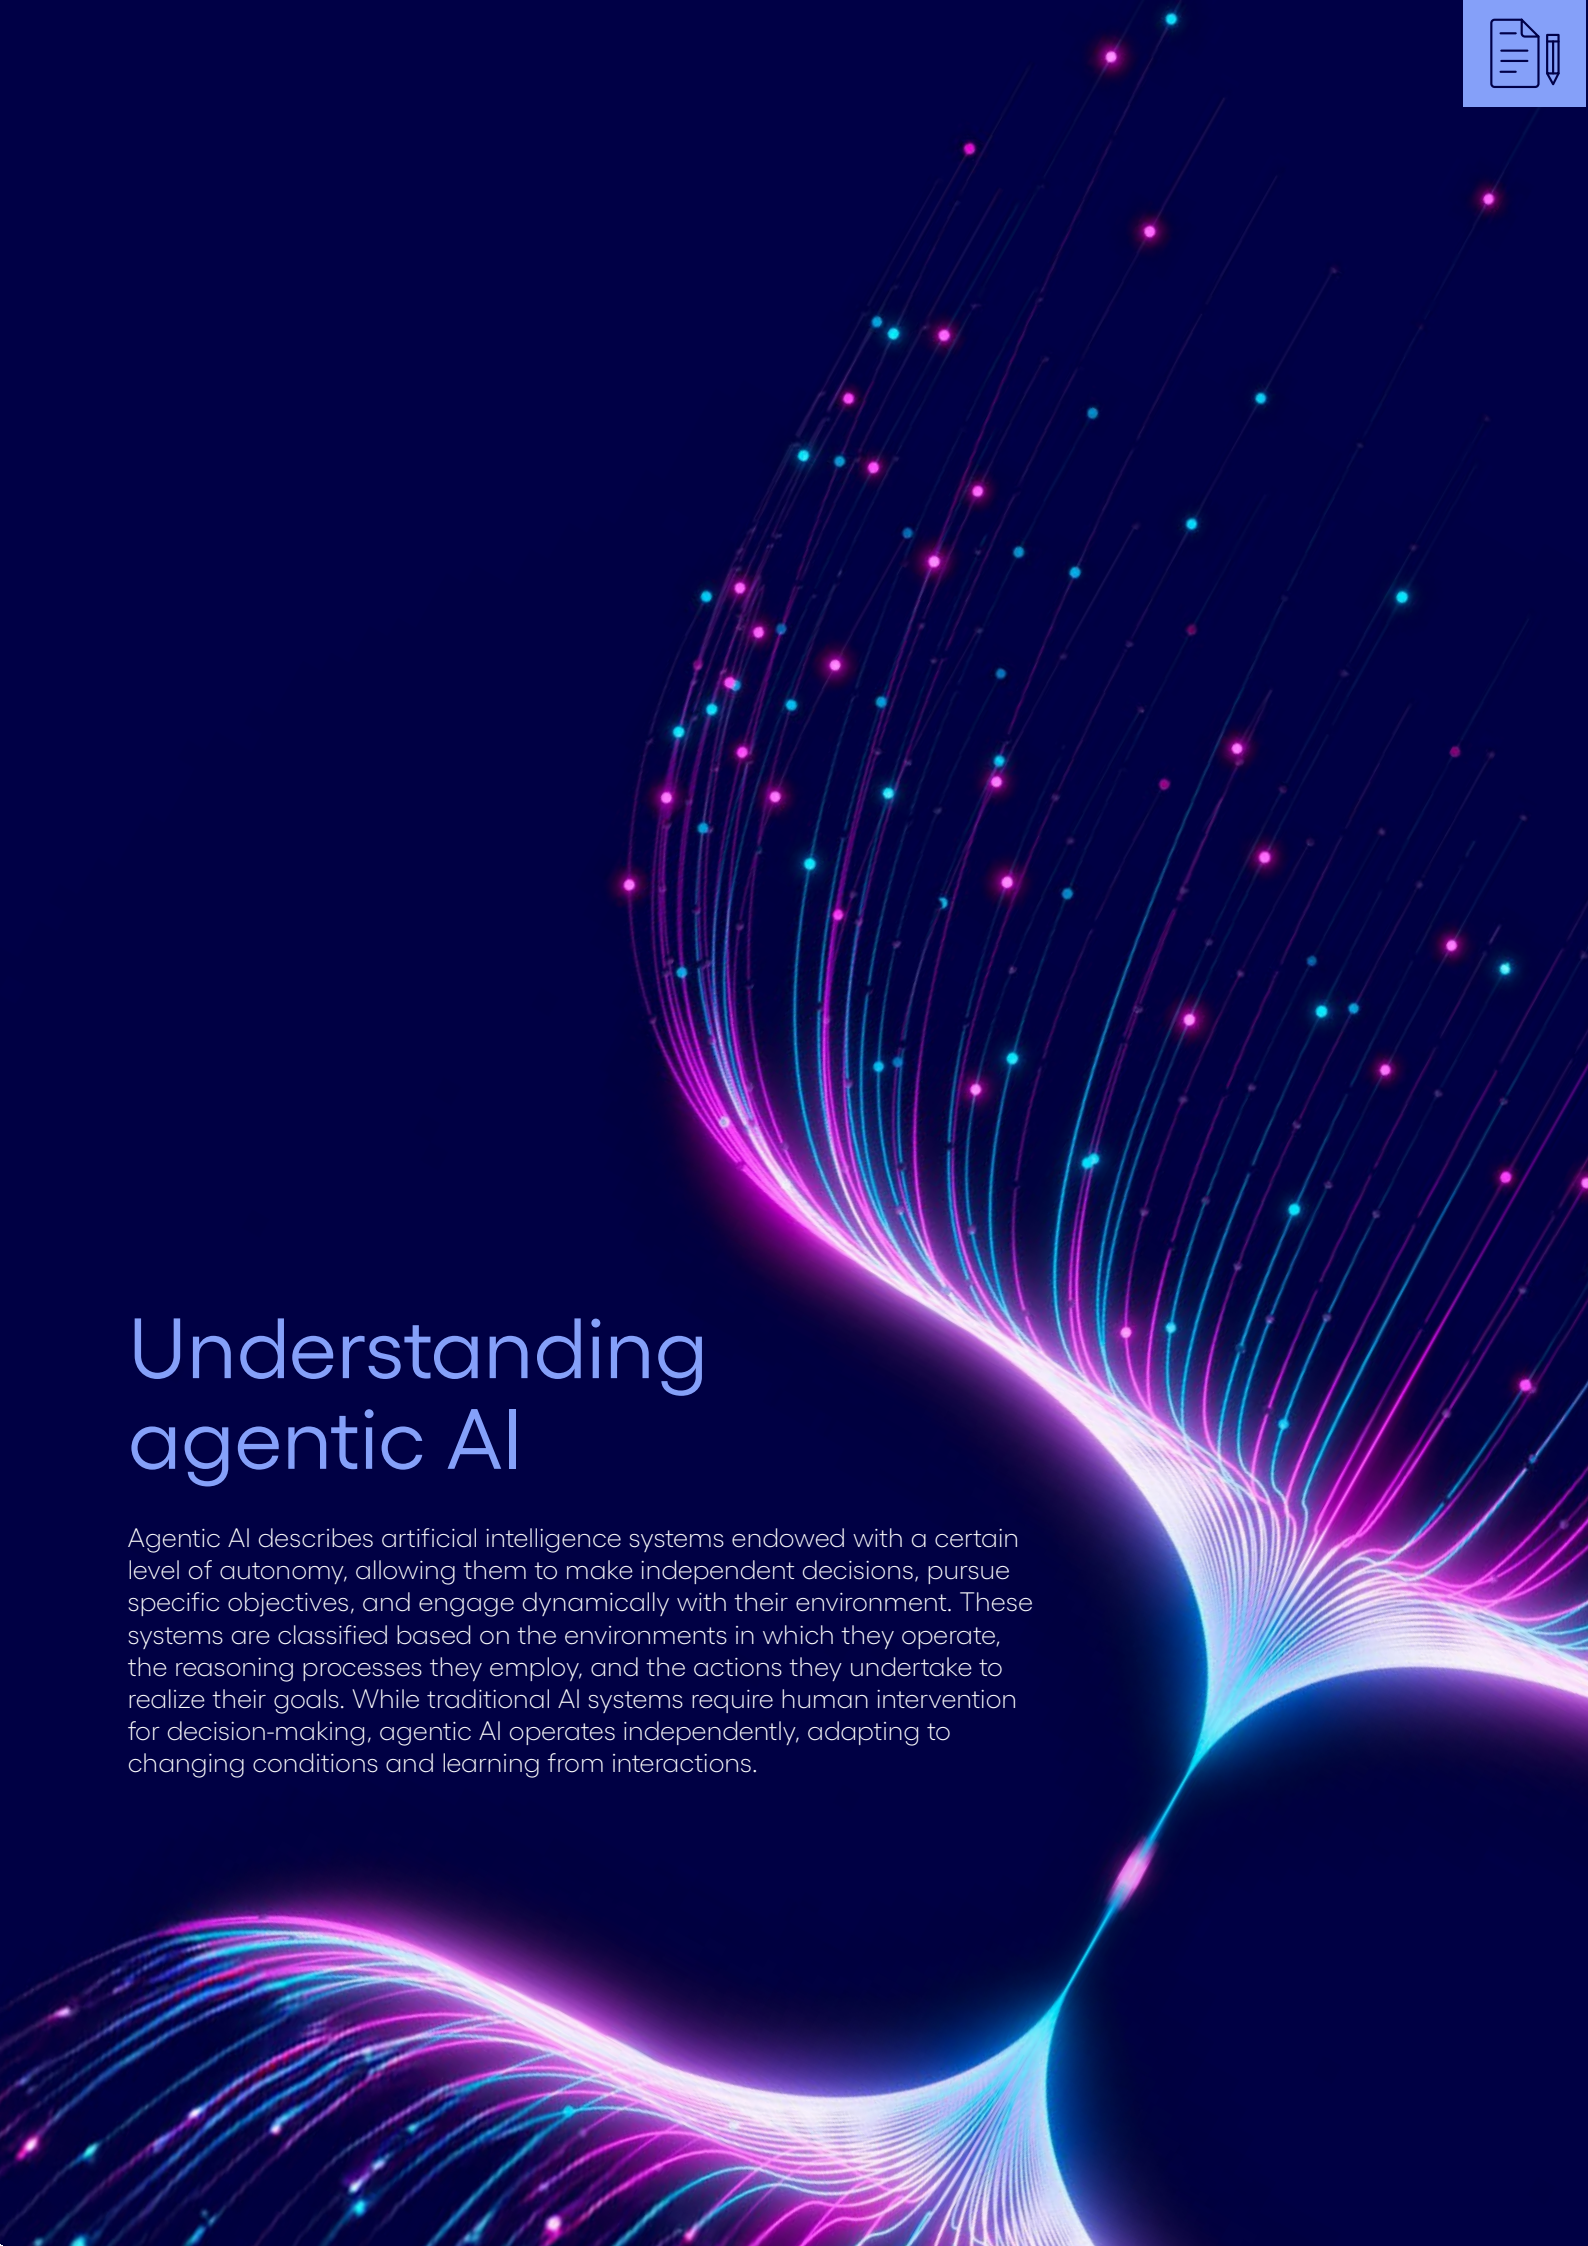

Understanding agentic AI

Agentic AI describes artificial intelligence systems endowed with a certain level of autonomy, allowing them to make independent decisions, pursue specific objectives, and engage dynamically with their environment. These systems are classified based on the environments in which they operate, the reasoning processes they employ, and the actions they undertake to realize their goals. While traditional AI systems require human intervention for decision-making, agentic AI operates independently, adapting to changing conditions and learning from interactions.

Core components of agentic AI

Perception
Understands input from the environment (text, images, etc.)

Cognition (LLM lives here)
Thinks, plans, reasons, and makes decisions

Action
Executes tasks or communicates results

Memory
Remembers past interactions and knowledge

Learning
Improves over time through feedback

Collaboration
Works with other agents or systems

Building blocks of agentic AI

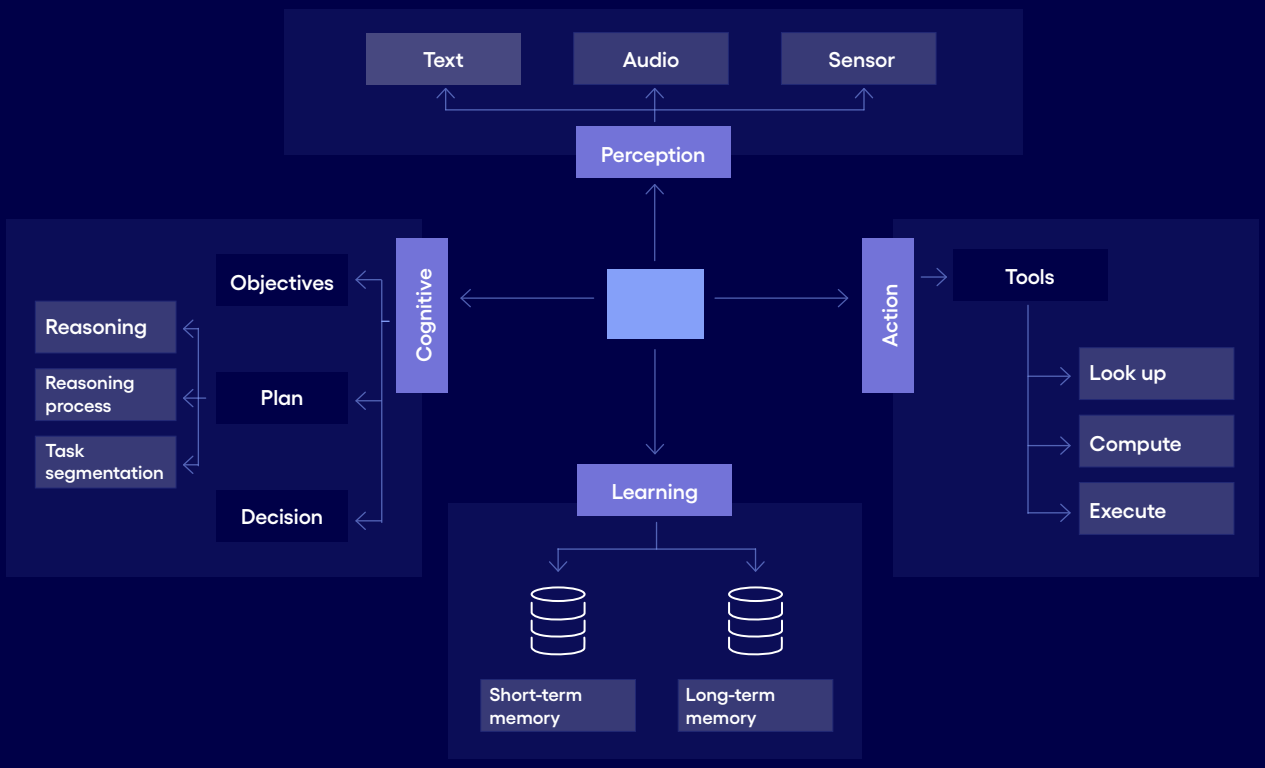

Technical frameworks for agentic AI

Key features of agentic AI frameworks

Integrated applications

Agentic frameworks feature built-in modules designed to facilitate intelligent automation, specifying both the behaviors of AI agents and the outcomes of their actions.

Planning and reasoning

Using a combination of **reasoning engines** adds intelligence to AI automation. These AI agents analyze information, create goals and plan several ways to complete them.

Monitoring and debugging

A robust and resilient design of agentic AI frameworks includes various monitoring and debugging tools that continuously track and enhance performance and eliminate bottlenecks.

Few popular frameworks - agentic AI

LangChain:
This is the most prevalent framework for building applications with large language models. It provides tools for connecting LLMs to other data sources and AI-powered apps, enabling more complex agentic workflows.

AutoGen:
This framework is developed by Microsoft for creating multi-agent AI applications. It enables communication and collaboration between multiple agents to solve complex tasks.

CrewAI:
This is an open-source framework that simplifies the orchestration of autonomous agents. It allows developers to create “crews” or teams of agents that can work together to achieve common goals.

LangGraph:
This framework leverages a graph-based methodology for constructing agentic AI workflows, making it especially well-suited for scenarios that involve intricate state management and the need to track dependencies.

Comparison of different agentic AI frameworks

Please see below diagram for comparison of different agentic AI frameworks with respect to their features, advantages and real-time applications.

Agentic AI Frameworks Comparison

	Features:	Advantages:	Applications:
ADK	Robust platform for advanced autonomous systems, integrates agent based AI solutions	Scalability, advanced efficiency, adaptability, interoperability, and real-time interactions	Demand forecasting, dynamic pricing, and inventory replenishment agents across stores. Travel disruption and hotel revenue management for optimization
LangGraph	Advanced NLP tool, performs sentiment analysis, text summarization and language understanding	Proactive decision making, superior data security and optimized information consumption	Personalized omnichannel shopping, conversational commerce journeys, multi-step travel planning and loyalty-driven recommendations
CrewAI	Integrates seamlessly with advanced team collaboration features, robust framework for projects	Predictive analytics, resource management, reporting and task automation	Collaborative agents for merchandising, promotions, and supply planning; guest experience and travel operations crews coordinating service and upsell
Microsoft Semantic Kernel	Powerful platform for data processing and analysis through semantic technologies	Understand and interpret human languages, sentiment analysis, entity recognition, knowledge graph	Enterprise-grade virtual store associates, policy-aware travel advisors and store agents
Microsoft AutoGen	Intelligent data processing; automation of routine tasks, scalability, enhanced security	Improved productivity, data accuracy, cost efficiency, reduced human error, compliance	Autonomous pricing and supply chain agents, booking optimization agents for travel sector
Smolagents	Improve operational efficiency, reliability & data accuracy, enhanced security and compliance	Provides accurate and reliable data, handles future growth, enhanced security compliance	Cost-effective automation for franchises, and budget hospitality
AutoGPT	Performance assessment with automated, objective metrics and fostering innovation	Memory management for context management, use plug-ins and execute code	Self-improving travel content, CX optimization, and growth strategy agents

Agentic AI Core - LLMs and RAG

Large language models (LLMs) are highly proficient in comprehending, processing and generating human language. Despite their capabilities, they are unable to access or incorporate real-time information from external sources. Their responses are constrained by the data they were originally trained on. As a result, this built-in knowledge gap can cause the information provided to be outdated or lacking in completeness.

Large language models (LLMs)

LLMs are advanced machine learning models trained on massive data sets of text and code. Their core capabilities are wide-ranging and include sophisticated text generation, accurate summarization of lengthy documents, seamless language translation, comprehensive question-answering, and even generating functional code.

Examples:

OpenAI GPT-4 | Meta AI Llama 3.1 | Claude

Key characteristics

1 **Training data**
LLMs are trained on substantial amounts of text from various sources, which helps in learning language models, syntax and contexts.

2 **Generate outputs**
LLMs can generate text that is intelligible and relevant to the input prompt.

3 **Useful systems**
LLMs are most commonly used for generating texts, summarizing data and handling language translation requests.

Applications of large language models

Conversational AI
Chatbots, virtual assistants, helpdesk automation, voice, multimodal and avatar agents

Customer support and chatbots
Helpdesk bots, virtual assistants, automated responses

Software development
Code generation, explanation and troubleshoot bugs

Education and tutoring
Personal tutors, language learning

Business intelligence and reports
Take insights from raw data and generate reports

Legal and compliance
Summarize legal documents, draft legal contracts, candidate screening

Academic research
Literature reviews, idea generation

Translation and multilingual support
Enhance communication, real-time language translation

Key benefits of LLMs

Scalability	LLMs can retrieve data from large datasets, making them suitable for large datasets.
Efficiency	LLMs can perform language related tasks very effectively such as translation, text generation, question answering etc.
Automation	LLMs can utilize intelligent processes across various functions such as customer service, data analysis and generate reports at the desired frequency level.
Improved user experience	LLMs can help in improving user experience with chatbots and search engines by assisting with more personalized responses.
Content generation	LLMs help in generating content at various levels such as text and code, making them useful for marketing and other creative fields.
Cross-functional integration	As per Forbes, LLMs are increasingly embedded across enterprise platforms (e.g., Google Workspace, Microsoft 365), enabling seamless integration into daily workflows and enhancing collaboration. **Source:** Forbes
Cost optimization	As per Gartner, with declining inference costs and hybrid deployment models (cloud + on-prem), LLMs offer scalable solutions with improved ROI and lower total cost of ownership. **Source:** Gartner
Sentiment and emotional analysis	As per Microsoft, LLMs can interpret tone and context to detect emotions and sentiment in text, helping improve brand monitoring and user engagement. **Source:** Microsoft

Limitations

Despite their impressive capabilities, pure large language models possess several inherent limitations that hinder their reliability and applicability in certain contexts, particularly within enterprise settings. One of the most significant challenges is their propensity for "hallucinations," where the models generate information that is factually incorrect, nonsensical, or entirely made up.

Another major constraint is the "knowledge cut-off." LLMs are trained on datasets up to a specific point in time, meaning their knowledge is limited to the information available during their last training update. For instance, the initial version of GPT-4 had a knowledge cut-off in September 2021, rendering it unable to provide information on events or developments occurring after that date. Furthermore, LLMs often suffer from a lack of transparency; it is difficult to trace the exact source or reasoning behind their generated answers, making verification a challenge.

Pure LLMs also struggle with domain specificity. They are not inherently equipped to handle proprietary, niche, or frequently updated enterprise data, which is crucial for many business operations. Attempting to update their knowledge base by retraining them for the latest information is computationally expensive and prohibitively slow, making it an impractical solution for dynamic data environments.

Large language models can inadvertently expose sensitive enterprise information, especially when handling proprietary or regulated data. "The Council of Europe" warns:

While pure LLMs exhibit these limitations, modern enterprise AI systems increasingly rely on tool augmentation, including retrieval-augmented generation (RAG), secure tool invocation, and controlled browsing; which significantly mitigates issues such as hallucinations, domain specificity, and knowledge cut-off. However, these extensions do not eliminate the risks entirely; they shift the challenge toward orchestrating reliable pipelines, validating retrieved data, securing tool interfaces, and ensuring auditability. Thus, understanding LLM limitations in isolation is necessary but insufficient when designing agentic AI systems.

Retrieval-augmented generation (RAG)

To address the inherent limitations of standalone large language models, particularly concerning factual accuracy, data freshness and domain specificity, the concept of retrieval-augmented generation (RAG) was introduced. RAG is not a new LLM, but rather a powerful framework that intelligently combines the vast generative capabilities of LLMs with external, up-to-date knowledge bases. Its primary purpose is to overcome issues like hallucinations and cut-offs by supplying the LLM with real-time, verifiable context before it generates a response.

The fundamental concept behind RAG is elegantly simple yet profoundly effective: Instead of relying solely on its internal, static training data, an LLM first retrieves relevant information from an external, dynamic data source. This retrieved information then acts as a contextual guide for the LLM's generation process, ensuring that the output is grounded in accurate and current facts. The RAG paradigm was first formalized and introduced by researchers at Facebook AI (now Meta AI) in a seminal paper published in 2020, marking a significant step forward in making LLMs more reliable and practical for real-world applications. Graph RAG uses graph-based retrieval to connect relationships between data points, enabling AI to deliver more structured and context-rich answers for complex business scenarios. Agentic RAG orchestrates multiple AI agents that collaborate to retrieve, synthesize, and reason over enterprise data, automating sophisticated workflows and supporting dynamic decision-making.

Source: AWS

Retrieval-augmented generation approach

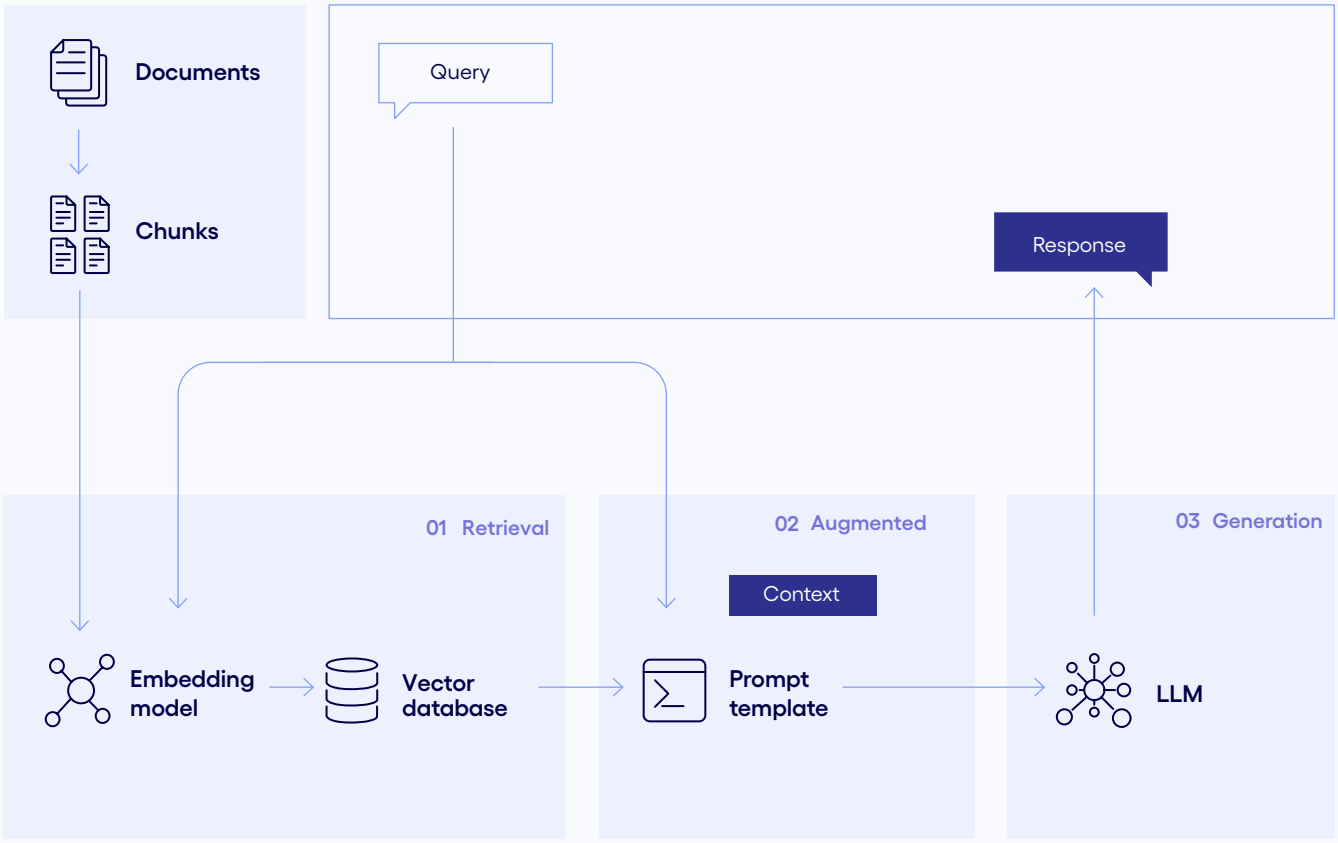

RAG developer's stack

LLMs

- Llama 3.3
- Phi-4
- Gemma 3
- Qwen 2.5
- Mistral
- DeepSeek
- OpenAI
- Claude
- Gemini
- Cohere
- Amazon

Frameworks

- Langchain
- Llama Index
- Haystack
- txtai

Vector database

- Chroma
- Pinecone
- Qdrant
- Weaviate
- Milvus

Data extraction

- Crawl4AI
- Firecrawl
- Scrape GraphAI
- Megaparser
- Docling
- Llama Parse
- Extract Thinker

Open LLMs access

- Hugging Face
- Ollama
- Groq
- Together AI

Text embeddings

- NOMIC
- SBERT
- BAAI BGE
- Ollama
- OpenAI
- Voyage AI
- Google
- Cohere

Evaluation

- Giskard
- Ragas
- TruLens

Transformative benefits of RAG implementations

The RAG provides a multitude of transformative benefits, fundamentally improving the utility and reliability of large language models across various applications. By grounding LLM responses in real-time, verifiable data, RAG can reduce the occurrence of hallucinations by an impressive percentage in factual question-answering scenarios, leading to much more trustworthy outputs.

RAG also addresses the critical issue of data freshness. It bypasses the inherent knowledge cut-offs of LLMs by incorporating the latest information from external knowledge bases, ensuring that responses are always current and relevant. This leads to increased transparency and trust, as RAG systems can often provide source attribution. From a practical standpoint, RAG offers substantial cost efficiency. It eliminates the need for expensive and time-consuming LLM fine-tuning or full retraining cycles whenever new data becomes available, making knowledge updates far more agile and economical.

RAG systems demonstrate strong domain adaptability. They excel in handling highly specific or proprietary enterprise contexts, such as legal documents, medical records or internal policy manuals, where general LLM knowledge would fall short. Ultimately, this leads to an improved user experience, as users receive more reliable, current and verifiable responses, enhancing their overall satisfaction and trust in AI applications.

Summary

Large language models and retrieval-augmented generation represent two powerful models in AI-driven text processing. While LLMs excel in creativity and adaptability, RAG stands out in accuracy and adaptability. Companies need to consider their specific requirements—whether they prioritize real-time accuracy or creative adaptability—when choosing between these two approaches. Together, LLMs and RAG pave the way for transformative advancements in how humans interact with technology, redefining the boundaries of possibility in natural language processing.

Agentic AI platforms

Agentic AI platforms are enterprise-grade systems designed to create, manage, and orchestrate autonomous AI agents that can perceive context, plan actions, and execute complex workflows with minimal human intervention. These platforms go beyond task automation to **deliver goal-driven outcomes,** transforming how organizations operate, innovate, and scale.

Key characteristics and capabilities

Autonomous AI agents
These platforms empower the creation of intelligent, software-driven tools (agents) that can perform tasks traditionally handled by humans.

Outcome-oriented
Agentic AI focuses on delivering real organizational value by automating complex workflows and achieving specific business outcomes, rather than just making employees more efficient.

Workflow orchestration
These platforms can orchestrate entire workflows, identify outcomes and automate actions, whether triggered by external events, system alerts, or internal requests.

Adaptability and decision-making
Agents are designed to be highly adaptive and can understand user needs based on business context, autonomously formulating plans and taking actions.

Integration and interoperability
Agentic AI systems can integrate with various systems and tools, allowing agents to access and manipulate data and execute actions across different platforms.

Governance and security
Introduce agentic governance frameworks to manage autonomy boundaries, data access, and escalation protocols. Ensure governance approach for agentic AI aligns with NIST's AI Risk Management Framework, grounding system design and oversight in the core functions of Govern, Map, Measure, and Manage to ensure trustworthy operational behavior. Structure these capabilities within an AI Management System consistent with ISO/IEC 42001, which establishes formal requirements for responsible AI development, risk assessment, and continuous improvement. Finally, ensure that the compliance roadmap incorporates the EU AI Act's phased obligations, including prohibitions and literacy requirements.

Source: McKinsey

Proactive behavior
Agentic AI takes initiative by recognizing patterns and anticipating what needs to be done next. It acts without waiting for instructions, often resolving issues or starting tasks before being asked—like a smart teammate that's always one step ahead.

Source: AWS

Examples of agentic AI platforms and frameworks

Krista AI
Pioneers agentic platforms with a focus on business outcomes, where agents orchestrate outcomes and perform actions autonomously.

Microsoft AutoGen
A platform designed to develop and deploy autonomous AI agents.

LangChain
A framework that facilitates the development of applications utilizing large language models, supporting the creation of agentic AI workflows.

CrewAI
A platform featuring multiple agents, engineered to optimize and simplify workflows in a wide range of industries.

Adept
Specializes in leveraging agentic AI to streamline software workflows, with the goal of automating various software development activities.

UiPath
Recognized as one of the leading agentic automation companies, specializing in robotic process automation (RPA) combined with AI capabilities.

Automation anywhere
Offers an AI platform that supports the simultaneous operation of multiple autonomous agents, each handling specific tasks while collaborating to achieve broader objectives.

Moveworks
A platform that focuses on AI-powered automation, including agentic AI capabilities.

Salesforce agentforce
An all-in-one agentic platform designed to empower enterprises to scale autonomous AI agents securely, leveraging low-code tools for seamless deployment and management.

NVIDIA AI platform
Enables efficient data management and access—an essential foundation for developing agile and responsive agentic AI applications.

Benefits of agentic AI platforms

Increased efficiency	Automating complex workflows and tasks can significantly improve efficiency and reduce human workload.
Reduced costs	Agentic AI platforms, with their expertise in automation capabilities, can help in reducing operational costs as the need for manual intervention is minimized.
Enhanced precision	Agentic AI systems are capable of executing tasks with superior accuracy and consistency compared to humans, leading to fewer errors and higher quality results.
Faster decision-making	Autonomous agents can make decisions and take actions quickly, enabling faster responses to changing business conditions.
Greater scalability	Agentic AI platforms are equipped to manage vast datasets and complex processes, empowering organizations to expand their operations and respond effectively to evolving business needs.
Better customer experience	Agentic AI can be used to personalize customer interactions and provide faster, more efficient service.

Applications of agentic AI and industry instances

Retail and Consumer Brands | Travel | Hospitality

- Agentic AI in the RCG industry
- Agentic AI in the travel industry
- Agentic AI in the hospitality industry

Agentic AI in the RCG industry

The industry has reached a point where the speed, volume, and complexity of decisions required to stay competitive exceed human capacity.

Below is a practical view of where agentic AI is making a real impact across the retail and consumer goods value chain.

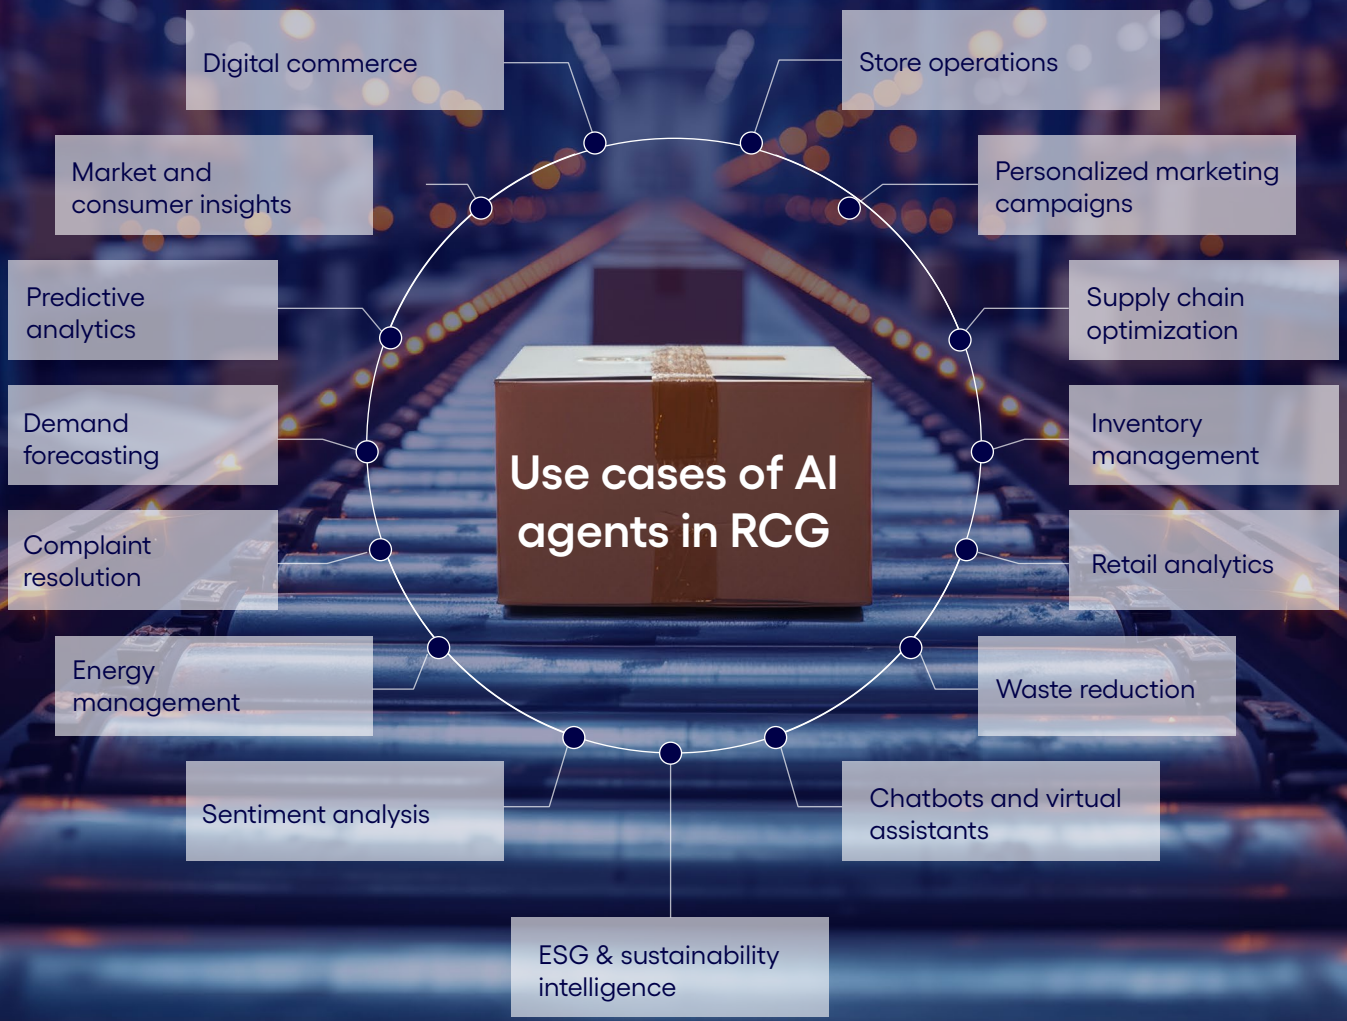

Enhanced operational efficiency

Automated tasks
Agentic AI can automate routine tasks like inventory management, order processing and demand forecasting, freeing up human employees for more strategic work.

Supply chain optimization
AI agents can analyze vast amounts of data to identify bottlenecks, optimize routes and predict disruptions in real-time, leading to a more efficient and resilient supply chain.

Personalized customer experiences

Predictive analytics
Agentic AI can analyze customer data to predict their needs and preferences, enabling businesses to offer personalized recommendations and tailored experiences.

Proactive customer support
AI agents can provide 24/7 customer support, addressing inquiries and resolving issues quickly and efficiently.

Hyper-personalized marketing
Agentic AI can help businesses create hyper-personalized marketing campaigns that resonate with individual customers, driving engagement and sales.

Agile and responsive businesses

Real-time decision-making
Agentic AI can make real-time decisions based on changing market conditions and customer behavior, enabling businesses to adapt quickly and stay ahead of the competition.

Up-to-date pricing
AI agents can analyze pricing data and adjust prices dynamically based on factors such as demand, competition and inventory levels.

Rapid product development
Agentic AI can help businesses identify new product opportunities and develop products that meet evolving customer needs.

Examples of agentic AI in retail and CPG

Inventory management
AI agents can optimize inventory levels, predict demand fluctuations and prevent stockouts.

Supply chain optimization
AI-driven agents oversee and enhance supply chain operations, proactively detecting possible disruptions and dynamically rerouting shipments in real time.

Customer service
Cognizant collaborated with a CPG client to develop a chatbot for call center agents to instantly retrieve relevant information from knowledge articles based on user queries, improving response time, accuracy, and overall customer service efficiency.

Marketing and sales
AI agents can personalize marketing campaigns, target specific customer segments and optimize sales processes.

ESG & sustainability intelligence
Agents track carbon footprint, packaging waste, and energy usage across the value chain, recommending actions to meet ESG targets.

Source: McKinsey

Key considerations for retail and CPG companies

Data quality
The effectiveness of agentic AI systems hinges on the availability of high-quality data, which is essential for generating accurate predictions and well-informed decisions.

Talent acquisition
Companies will need to invest in talent with expertise in AI and data science.

Ethical considerations
It's important to ensure that agentic AI systems are used ethically and responsibly.

Collaboration
Agentic AI systems should be designed to work in collaboration with human employees, rather than replacing them.

Agentic AI in the travel industry

In travel, every delay, rebooking, or lost connection triggers a cascade of decisions that must happen instantly to keep customers moving. With disruptions now more frequent and traveler expectations rising, the industry needs systems that can anticipate problems and act before they impact the journey. Agentic AI fills this gap—coordinating decisions and actions in real time to reduce friction, personalize service, and stabilize operations.

Here are the real-world travel moments where agentic AI is already making a measurable difference.

Hyper-personalized travel experiences

- Agentic AI can analyze individual customer data (preferences, booking history, etc.) to create tailored travel recommendations and itineraries.
- This includes suggesting activities, restaurants and accommodation based on individual interests and preferences.
- Depending on real-time conditions such as flight delays, itineraries can be adjusted dynamically.

Efficiency and automation

- Agentic AI can automate tasks that previously required human intervention, such as check-ins, booking amendments and customer support, resulting in expedited response times, reduced operational costs and improved customer satisfaction.
- AI agents can handle routine tasks, freeing up human travel agents to focus on more complex issues and personalized service.

Real-time problem-solving

- When disruptions occur (e.g., flight cancellations, weather delays), agentic AI can immediately take action to rebook flights, notify travelers and find alternative routes.
- For example, Cognizant helped develop AI-powered virtual agents for Travel industry clients to handle policy-related customer service queries. These agents will ensure faster response times, consistent support, reduced manual load, and improved customer satisfaction through intelligent automation and natural language understanding.

Loyalty program optimization

- AI agents analyze traveler behavior and spending to personalize loyalty rewards and tier upgrades.
  Source: McKinsey

Revenue optimization

- Agentic AI analyzes booking trends, no-show rates and external factors to fine-tune overbooking strategies and maximize seat occupancy. This intelligent automation boosts revenue while minimizing passenger disruptions.
  Source: McKinsey

Proven agentic AI use cases above and below the wing

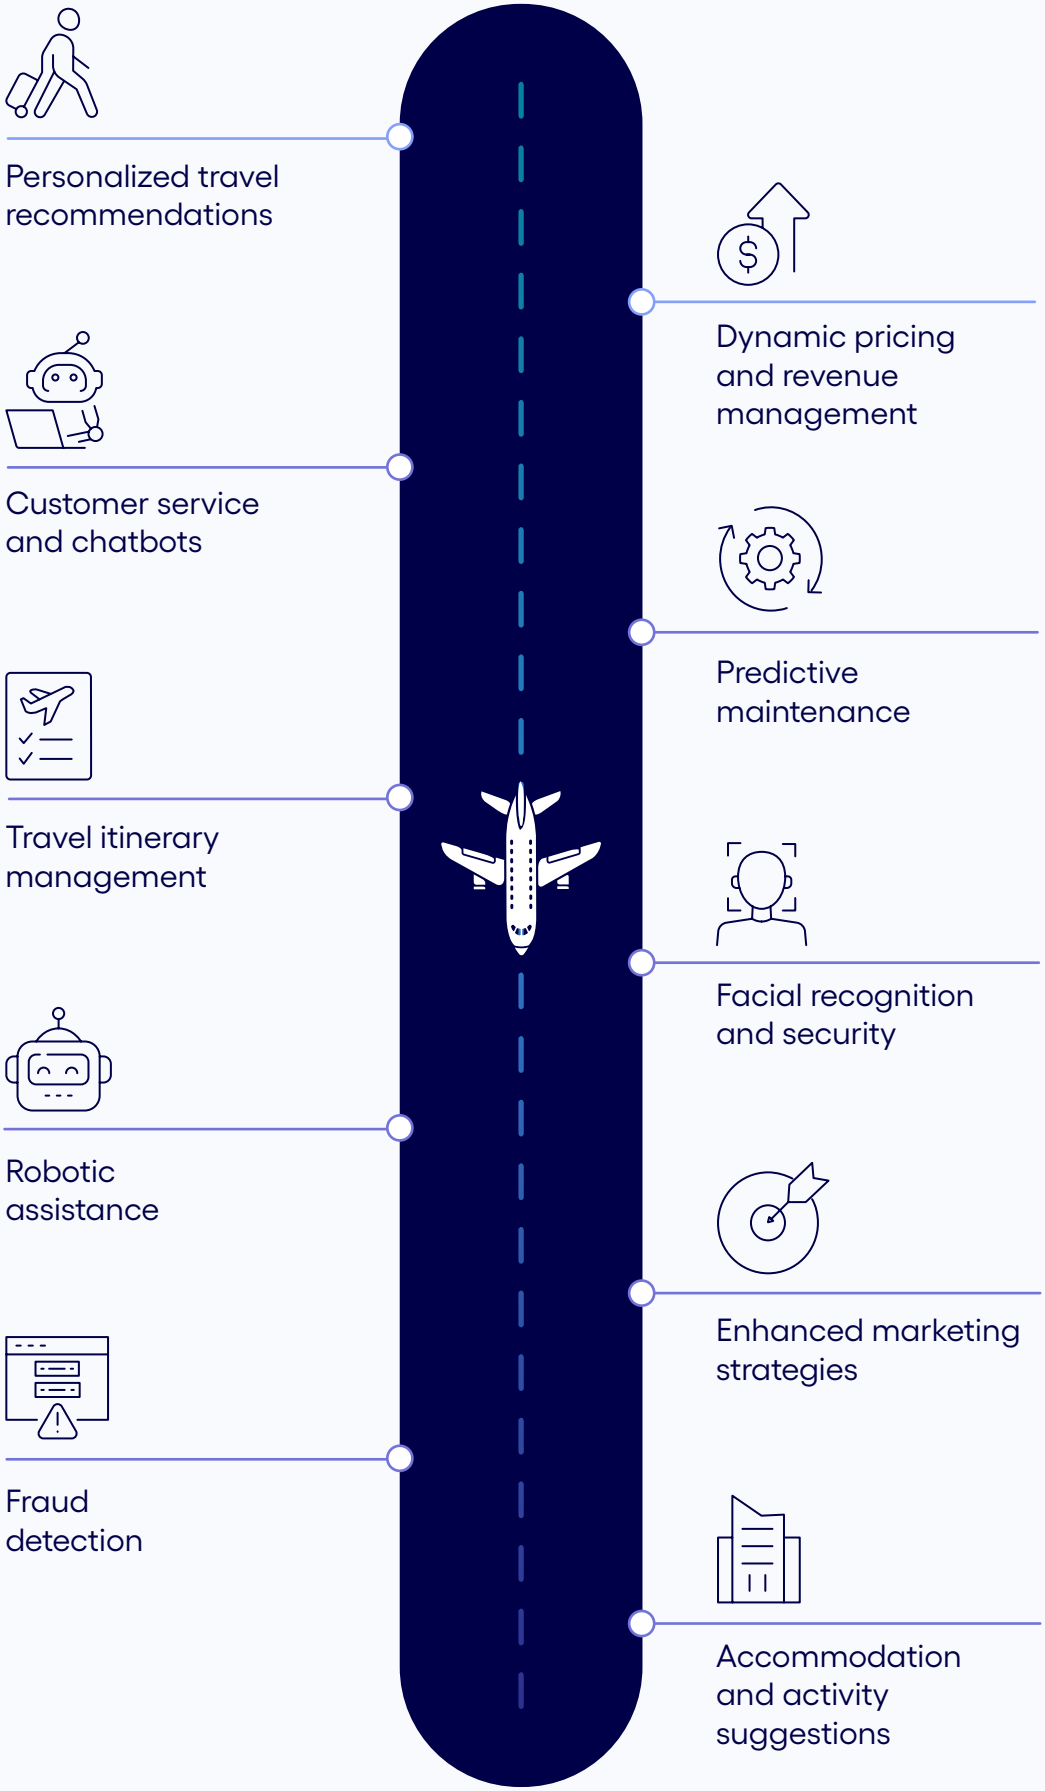

Agentic AI in the hospitality industry

Agentic AI systems are sophisticated artificial intelligence programs characterized by their ability to act autonomously, make decisions and pursue predefined goals without constant human intervention.

The applications of agentic AI in hospitality are vast and varied. They can span the entire guest journey, from optimizing seamless booking processes by predicting preferred room types and dates, to delivering highly personalized in-stay services such as adjusting room climate based on guest habits, or even autonomously arranging transportation and dining experiences tailored to individual preferences. These systems aim to create an intuitive, anticipatory service environment.

Major challenges in the hospitality industry

Challenge 1
Data privacy, security and trust

The implementation of agentic AI, while offering immense benefits, introduces significant challenges concerning data privacy, cybersecurity and the fundamental trust guests place in hospitality providers. The nature of autonomous systems means they interact with, collect and process vast amounts of highly sensitive personal and behavioral data.

Sensitive data handling
To provide hyper-personalized experiences, agentic AI systems need access to extensive guest profiles. This includes not only basic demographic information but also preferences for room temperature, dietary restrictions, communication styles, past purchase history, preferred activities and even movement patterns within the property via IoT devices. Managing such a rich and sensitive dataset requires robust data governance frameworks, strict access controls and transparent policies to ensure guest privacy is paramount. Any mishandling or breach could have severe reputational and financial consequences.

Cybersecurity vulnerabilities
As agentic AI systems become more interconnected with various hotel operational systems (PMS, POS, smart room controls, maintenance systems), the overall attack surface for cyber threats expands. Autonomous systems, if compromised, could potentially give attackers deep access to sensitive guest data and critical operational controls. This necessitates world-class cybersecurity infrastructure, continuous monitoring and proactive threat detection capabilities to protect against breaches, ransomware and other malicious activities.

Gaining guest loyalty
Hotels must establish a strong assurance of ethical AI use and data protection to lighten fears and encourage adoption of AI-enhanced services.

Challenge 2
Impact on the workforce

While agentic AI excels at efficiency and personalization, it faces significant hurdles in replicating the nuanced “human touch” that often defines exceptional hospitality. Furthermore, its integration raises important questions about its impact on the human workforce.

Service nuance
Hospitality thrives on empathy, spontaneous problem-solving and genuine human connection. Agentic AI, despite its advanced capabilities, struggles to replicate the subtle cues of human interaction, such as understanding unspoken emotions, offering truly empathetic responses in challenging situations, or engaging in charming, unscripted conversations. For complex, emotional or highly personalized issues, guests overwhelmingly prefer human interaction. As such, AI should augment, rather than fully replace, human staff in high-touch guest service roles.

Job displacement concerns
The automation capabilities of agentic AI, particularly for routine and repetitive tasks (e.g., check-in, basic concierge services, scheduling), can lead to concerns about job displacement among frontline hotel staff. This apprehension needs to be managed proactively through clear communication and strategic workforce planning, focusing on how AI will change, rather than eliminate roles.

Reskilling imperative
Instead of outright displacement, the greater impact will be a transformation of job roles. Staff will need to learn how to monitor and manage AI systems, interpret AI-generated insights, and focus on higher-value tasks that require emotional intelligence, creative problem-solving and complex decision-making. This shift requires comprehensive training programs to ensure employees are equipped for this collaborative human-AI environment.

Here's where AI agents are quietly transforming hospitality behind the scenes.

- A leading player in India's hospitality sector utilizes an AI-powered dynamic pricing engine to modify room rates in real time, considering variables such as historical booking patterns, local happenings and competitor prices.
- Cognizant partnered with an American hospitality company to enable a seamless customer experience from planning to booking to post-stay across client's multi-brand ecosystem. Leverage integrated data, AI, and personalization to enhance satisfaction, drive loyalty, and deliver consistent service throughout the entire guest lifecycle.
- Restaurants: In the restaurant industry, AI enables more efficient menu optimization, anticipates customer demand and modifies recommendations to individual preferences. For example, Cognizant partnered with a major pizza chain to launch a scalable voice assistant to improve customer experience, and establish strong technology base for future AI-driven customer support and ordering automation.
- Supply chain management: AI can help with tracking shipments, and identify any delays in procuring goods.
- A global hospitality leader recognized for its AI-driven sustainability efforts has adopted intelligent analytics to minimize food waste across its global operations. By analyzing real-time data on food production and disposal, these tools support chefs in refining portion sizes, cooking techniques and inventory decisions—leading to notable reductions in waste and emissions.
  Source: Abode Worldwide

Advantages of agentic AI in the hospitality industry

Advantage 1
Hyper-personalized guest experiences

Tailored recommendations
Agentic AI can analyze vast amounts of guest data—past stays, preferences, social media activity and real-time behavior—to offer tailored suggestions for food and beverage, local activities, and even in-room amenities. This level of customization moves beyond basic preferences to anticipate implicit desires.

Proactive service
This might include pre-booking airport transportation based on their flight schedule or even dispatching a maintenance agent to a room for early signs of malfunctioning refrigerator, all without the guest needing to make a request.

Enhanced communication
AI agents can manage personalized communication flows, from sending tailored pre-arrival guides with local tips and essential check-in information, to crafting post-stay follow-ups that include personalized offers for future visits or solicit specific feedback. This continuous, intelligent communication ensures guests feel valued and understood throughout their entire journey, fostering deeper engagement and loyalty.

Autonomous guest acquisition & booking
Agentic AI agents can autonomously monitor guest calendars, preferences, and loyalty data to proactively book rooms, reserve amenities, and even arrange transportation—without guest input.
Source: McKinsey

Multi-language & cultural personalization
AI agents communicate in guests' native languages and adapt recommendations based on cultural preferences.

Impact

The cumulative effect of these hyper-personalized services is a dramatic increase in guest satisfaction. Studies suggest a potential increase in guest satisfaction, directly correlating with higher rates of repeat bookings and positive word-of-mouth referrals. By making each guest feel uniquely cared for, hotels can build stronger, more lasting relationships.

Advantage 2
Streamlined operations and efficiency

Agentic AI offers significant advancements in operational efficiency by automating and optimizing various back-of-house tasks. This leads to substantial resource optimization and cost savings across the board.

Automated task management
AI agents can dynamically optimize complex schedules, such as housekeeping assignments based on real-time check-outs and arrivals, guest preferences for cleaning times, and staff availability. They can also manage inventory levels for supplies, food and beverages, automatically reordering when thresholds are met. Furthermore, routine maintenance tasks can be scheduled and dispatched autonomously, ensuring facilities are always in optimal condition.

Predictive maintenance
Agentic AI systems can monitor equipment performance (e.g., HVAC systems, elevators, kitchen appliances) in real time using IoT sensors. By analyzing patterns and anomalies, they can identify potential equipment failures before they occur, triggering maintenance alerts or even ordering parts autonomously.

Effective resource utilization
AI agents can allocate resources effectively and they can analyze booking forecasts, event schedules and historical data to predict fluctuations in guest demand. Based on these predictions, the AI can make intelligent adjustments to staffing levels for different departments (front desk, F&B, housekeeping), ensuring optimal staff utilization and reducing unnecessary labor costs. This dynamic adjustment capacity minimizes both overstaffing and understaffing, leading to smoother operations.

Smart energy management
Agentic AI enables hotels to reduce energy consumption by autonomously managing smart systems such as lighting and thermostats based on environmental conditions, room occupancy, and guest behavior. This intelligent automation minimizes energy waste, lowers operational costs, and supports sustainability goals.

Advantage 3
Dynamic revenue optimization

Agentic AI provides an incomparable ability to maximize revenue by moving beyond traditional static pricing and offering highly personalized opportunities. This is critical in the hospitality landscape.

Dynamic pricing
Agentic AI can process vast quantities of real-time data, including competitor rates, local event schedules, weather forecasts, flight arrival patterns and even social media sentiment. It then dynamically adjusts room rates and package prices in milliseconds to ensure optimal occupancy and maximum revenue per available room (RevPAR). This means prices can fluctuate several times a day to capture peak demand or stimulate bookings during quieter periods, responding instantly to market shifts.

Tailored cross-selling
AI agents can identify the ideal moment and the most suitable offer for each individual guest. For example, if a guest has a high-loyalty status and has previously shown interest in spa services, the AI might offer a discounted spa package at check-in. These tailored suggestions are far more effective than generic offers, significantly increasing conversion rates for additional services and higher-tier rooms.

Forecasting accuracy
The ability of agentic AI to ingest and analyze diverse datasets allows for a significant improvement in demand prediction. By integrating internal data with external factors like local holidays, concert schedules, and even flight delays, AI can predict future demand with up to 20% greater accuracy. This enhanced forecasting allows hotels to make smarter decisions regarding inventory management, staffing levels and marketing campaigns, minimizing lost revenue from unsold rooms or under-optimized pricing.

The journey ahead

As organizations shift from small pilots to enterprise-scale agentic AI adoption, one theme has become unmistakable: true advantage will belong to those who can operationalize autonomous agents reliably, securely, and at scale. Insight alone is not enough. What enterprises need next is a repeatable system for designing, deploying, governing, and improving agentic solutions across their business.

Cognizant's AI Builder strategy—anchored by Cognizant Agent Foundry and Neuro® AI multi-agent orchestration—offers exactly that path forward. It provides a structured, end-to-end approach that enables organizations to move confidently from experimentation to industrialized, agent-driven transformation.

A modern framework for designing enterprise-grade agents

Cognizant Agent Foundry supplies standardized components, domain-specific small language models, reusable agent templates, and deep partner ecosystem integrations, giving enterprises a safe, scalable foundation for building agents aligned to real business goals.

A multi-agent architecture built for real-world complexity

Powered by Neuro® AI, Cognizant's multi-agent orchestration enables agents to collaborate across workflows, coordinate decisions, and manage task handoffs with embedded governance and human-in-the-loop controls, essential for operational environments where precision and oversight matter.

A proven methodology to discover, design, build & scale agentic solutions

Through the Agent Foundry lifecycle, organizations gain a repeatable, transparent path from use-case discovery to scaled deployment, grounded in feasibility assessments, business-value modeling, risk and compliance frameworks, and enterprise change management.

Conclusion

As the industry transitions from incremental automation to truly agentic operations, Cognizant provides the building blocks to turn vision into execution. With Cognizant Agent Foundry and our AI Builder approach, enterprises can move beyond isolated pilots and begin architecting an agent-powered future, one where autonomous systems work alongside people, orchestrate complex workflows, elevate decision quality, and accelerate outcomes across the value chain. The next step is not simply adopting agentic AI, but industrializing it, and Cognizant is uniquely positioned to guide that journey.

Authors

Anup Prasad

Senior Vice President, Consumer Business Unit Head

Retail, CPG, Travel & Hospitality

Anup Prasad is SVP & Business Unit Head for the Consumer Business (Retail, CPG, Travel & Hospitality) in Americas. In his role as BU Head, Anup is responsible for **strategy, growth & operational performance** of Cognizant's Consumer Business that serves some of the world's most recognized brands in retail, consumer goods, travel, restaurant & hospitality sectors.

Anup has been with Cognizant for ~25 years and has been responsible for winning and nurturing many marquee client relationships. He thrives at the confluence of technology & industry expertise and is known for his client-centricity & building high-performance teams. He has been recognized as one of the "Top 25 Consumer & Retail Consultants & Leaders".

Prior to his current role, Anup was heading CPG and Travel & Hospitality sectors across Americas and before that retail and consumer goods business for the Central US region at Cognizant. He started his career at Cognizant as a business analyst and has spent all his time in the Retail & Consumer industries growing up through the ranks.

Before beginning his career at Cognizant in 2001, Anup worked as an SAP Consultant at IBM and was part of core team implementing SAP at Arvind Mills in India. He earned a Bachelor of Technology from the Indian Institute of Technology, Delhi and an MBA from XLRI – Xavier School of Management, Jamshedpur in India. Anup relocated to US in 2001 after joining Cognizant and calls Dallas, TX his home.

LinkedIn

Vinay Swarup

AVP – Strategy & Business Development, Consumer Business

Retail, CPG, Travel & Hospitality

Vinay leads the strategy and business development for the Consumer Business Unit, backed by close to 21+ years of consulting expertise in the Retail, CPG and Travel & Hospitality. As a trusted advisor to CIOs and CDAOs, he specializes in developing data and AI strategies, bringing valuable insights for creating effective go-to-market strategies for Consumer Business clients.

Before joining Cognizant, Vinay managed numerous digital and data transformation projects for prominent Retail, CPG and airline companies. His experience includes setting up and leading data and AI CoEs at top global consulting firms, demonstrating his capability in spearheading innovation and growth in digital transformation and data analytics.

cognizant®

Cognizant helps engineer modern businesses by helping to modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. To see how Cognizant is improving everyday life, visit them at **www.cognizant.com** or across their socials @cognizant.

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

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

WF4464250