VERSION NEXT

AGENTIC AI ISSUE | SPRING 2025

AUTONOMOUS DECISION-MAKING

Advancing Problem-Solving in Al

Autonomous Decision-Making: Advancing Problem-Solving in Al

Enterprise operations face constant pressure to improve efficiency, adaptability and innovative capacity. Traditional rule-based AI systems, despite their initial promise, have consistently underperformed in delivering comprehensive solutions that can adapt to the dynamic nature of modern business environments. Their strict adherence to predetermined rules fundamentally limits their ability to learn, adapt and make autonomous decisions in response to changing circumstances. Agentic AI addresses these limitations by introducing systems that can continuously evolve their approaches based on new information and experiences, creating a paradigm shift in how organizations can leverage artificial intelligence for business success.

This emerging technology employs multiple AI agents that leverage large language models (LLMs) to enhance decision-making and understanding. The distinguishing feature of agentic AI lies in its ability to not only transform data into knowledge but make decisions and perform actions. This advancement enables end-to-end enterprise solutions to achieve business outcomes across diverse industries, from customer service optimization to scientific discovery and environmental sustainability applications. As with any new system, implementation can face some hiccups. While cloud service providers offer robust tools for adopting these systems, organizations face numerous challenges during development and deployment. Early adopters, however, who implement effectively while managing risks appropriately will likely establish sustainable advantages that prove difficult for competitors to overcome through later adoption efforts.

2 TEKsystems.com | VERSION NEXT, NOW



FEATURES

THE CHANGE AGENT

Our TEKsystems Global Services (TGS) experts explain that technology investments—especially in Al—are top priorities for businesses right now to improve efficiency and problem-solving. Agentic Al, the latest generation of Al, offers the groundbreaking ability to translate knowledge into autonomous action to facilitate productivity and innovation.

PAGE 7 ≫

MARKET PERSPECTIVE

Nader Mherabi, EVP, vice dean, and chief digital and information officer at NYU Langone Health, and David Brauchler, technical director at NCC Group, share how businesses can leverage agentic Al to boost business operations with effective, adaptable solutions across industries.

PAGE 13 ≫



03

OUR PERSPECTIVE

Organizations can balance human creativity with AI efficiency and data access by implementing agentic AI for better decision-making, cost savings and competitive advantage.

PAGE 19 ≫



Technology investments—especially in Al—remain top priorities for businesses right now to improve efficiency and problem-solving. Agentic Al, the latest generation of Al for enterprises, offers the groundbreaking ability to translate knowledge into autonomous action to facilitate productivity and innovation.

THE CHANGE AGENT

Reimagining Problem-Solving With Human-Machine Collaboration

Agentic AI operates through autonomous AI agents specifically designed to perform complex tasks by interpreting contextual information, making decisions based on that interpretation and executing actions aligned with predetermined objectives. This functionality bridges the traditional gap between static programming approaches and the dynamic adaptability required for modern enterprise challenges. By combining sophisticated decision-making capabilities with contextual understanding and adaptive behavior, agentic AI enables the automation of intricate workflows that previously required substantial human oversight and intervention.

The operational foundation of agentic AI rests on several key components working in coordinated harmony. Decision-making algorithms represent a central element, allowing the system to evaluate multiple options and select appropriate actions based on goals, constraints and contextual factors. These algorithms typically employ probability assessments, pattern recognition and objective alignment to make choices that optimize outcomes within given parameters. This capability enables agentic AI to navigate complex decision landscapes autonomously, making judgment calls that previously required human expertise.

Reinforcement learning mechanisms provide agentic AI with the ability to improve through experience, gradually enhancing performance by understanding which actions lead to desired outcomes under various conditions. This learning approach allows these systems to evolve and refine their strategies over time without explicit reprogramming. Essentially, they "learn by doing" through repeated interactions with their operational environment. The result is continuous improvement in performance and decision quality as the system accumulates experiential data and refines its understanding of effective approaches.

Contextual analysis capabilities enable agentic AI to interpret the nuances of situations, understanding not just explicit instructions but implicit context that informs more appropriate responses. This sophisticated level of comprehension allows the system to grasp stated requirements and unstated expectations, leading to more intuitive and appropriate actions across diverse scenarios. When combined with adaptive behavior mechanisms that modify strategies based on new information or changing circumstances, this contextual understanding ensures agentic AI remains effective even in dynamic and evolving environments.

More sophisticated implementations often incorporate multi-agent coordination architectures, where multiple specialized AI agents collaborate to achieve broader objectives. In these systems, individual agents focus on specific aspects of a task while coordinating their efforts through structured communication and task allocation frameworks. This distributed approach enables more complex problem-solving by combining specialized capabilities into coherent, coordinated action toward shared goals. Despite their dependence on initial programming and high-quality training data to develop their decision-making capabilities effectively, organizations will see increasing benefits in leveraging the remarkable autonomy and adaptability of agentic AI as a driver of business modernization.





Multiple industry forecasts, corroborated by extensive research across sectors, predict that AI agents will automate up to 70% of office work tasks within the next decade.¹

CHAR

(12M

Core Components and Architecture

A TYPICAL MULTI-AGENT SYSTEM ARCHITECTURE CONSISTS OF:



1. Specialized agents: individual AI agents with specific functions and capabilities



2. Shared memory: a repository for communication and knowledge sharing between agents



3. Orchestration layer: coordinates agent activities and workflow



4. Data storage and retrieval layer: manages information access and storage



5. Service layer: delivers AI capabilities across platforms





EXPERT SYSTEMS

Al advances toward agentic decisionmaking with MYCIN, which demonstrates its potential in medical diagnosis.



DEEP LEARNING TAKES OFF

Breakthroughs in deep learning drive major advances in image and speech recognition, kicking off a new era of autonomous Al.

2020

ETHICAL USE OF AI

As AI becomes more agentic, it raises questions about ethical use, accountability and surrounding regulations.



History of Agentic Al

DARTMOUTH CONFERENCE

The birth of AI as a field: John McCarthy and Marvin Minsk explore how machines could mimic human intelligence.

1980

DEEP BLUE VS. KASPAROV

IBM's supercomputer beats the world chess champion, proving AI can handle strategy and complex decision-making.

RISE OF LLMS

2025

2012

Al handles more complex conversation and content generation, narrowing the gap between human and machine.

Fig. 1. "History of Agentic Al"²



MARKET PERSPECTIVE

Nader Mherabi, EVP, vice dean, and chief digital and information officer at NYU Langone Health, and David Brauchler, technical director at NCC Group, <u>share how businesses can leverage agentic Al</u> to boost business operations with effective, adaptable solutions across industries.

02

Agentic Al Considerations From NYU Langone Health and NCC Group

How should companies implement agentic AI?

Nader Mherabi (NYU Langone Health): When companies start deploying AI, they should put together a strategic plan to make sure that the right technology is being applied to the right use case. You can't have AI outside of the workflow because people won't use it.

14 TEKsystems.com | VERSION NEXT, NOW

VERSION NEXT, NOW | TEKsystems.com 15

- 1



Will data be secure?

David Brauchler (NCC Group): Companies should recognize potential security risks associated with agentic AI. Everything that you add to your application environments or networks is additional attack networks that you have to govern. Security testing should be conducted when considering adding AI to systems and platforms lacking design that includes AI or AI-related risk in mind.









Nader Mherabi (NYU Langone Health): Agentic AI is not a set-it-and-forget-it type of application. It needs to be continuously monitored to ensure that the data is accurate, is being used appropriately and benefits end users.



OUR PERSPECTIVE

Organizations can balance human creativity with Al efficiency and data access by implementing autonomous Al agents for better decisionmaking, cost savings and competitive advantage.



Agentic Al Business Impact and Applications

The practical applications of agentic AI already demonstrate transformative value across many business functions and industries, optimizing processes that previously required substantial human oversight and intervention. By combining autonomous decision-making with contextual understanding and adaptive behavior, these systems enable organizations to achieve unprecedented levels of efficiency, responsiveness and innovation at scale. From customer service enhancements to operational optimizations and scientific breakthroughs, agentic AI for enterprises is reshaping how organizations approach complex challenges and opportunities.

ETHICAL CONSIDERATIONS AND ACCOUNTABILITY

As enterprise AI solutions gain greater decision-making autonomy, determining responsibility when an autonomous system makes errors or produces unintended consequences becomes increasingly complex, highlighting the importance of clear accountability frameworks across the business, such as the National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF).

Aligning AI systems with human values and ethical principles is critical during implementation. This will prevent actions that conflict with organizational values or broader societal norms, particularly in complex or ambiguous situations where multiple competing priorities exist. Explicit value alignment processes, ethical review mechanisms and continuous monitoring for value drift are essential elements of responsible agentic Al implementation.

ENVIRONMENTAL IMPACT CONCERNS

The environmental footprint of AI systems requires serious review during implementation planning and ongoing operations. The proliferation of data centers housing Al infrastructure contributes significantly to hazardous electronic waste generation. Implementation must incorporate sustainable life cycle management from deployment through eventual decommissioning to minimize environmental impact.

Al infrastructure development and operation depend heavily on critical minerals and rare earth elements. These components are frequently extracted through mining practices with substantial environmental and social consequences as the elements often appear in electronic waste streams. Although proper recycling and resource recovery programs aren't available, some strategies could reduce e-waste generation by up to 86%, including using equipment for longer (lifespan extension), refurbishing and reusing components, and designing hardware to facilitate recycling and upgrades.

Data centers powering AI infrastructure also utilize massive electricity resources that contribute to greenhouse gas emissions, depending on energy sources. Similarly, water consumption for cooling systems can further strip scarce resources in many regions, necessitating strategies like renewable energy adoption, efficiency improvements and water conservation measures for sustainable technology implementation.

TECHNICAL AND OPERATIONAL CHALLENGES OF AGENTIC AI



1. Agent Communication and Orchestration Issues

One of the most significant challenges involves creating effective communication channels between agents. Community forums highlight a specific technical issue where developers struggle with multidirectional communication: "I can route from A to B, but then B cannot go back to A after it has completed."

This particularly affects workflows that require bidirectional communication between parent and child agents. The system often generates "loop detected" errors when attempting to implement such communication patterns.



2. Scalability and Performance Bottlenecks

As multi-agent systems grow in complexity, scalability becomes a critical concern. While individual agents might operate efficiently, orchestrating numerous agents across an enterprise introduces significant performance challenges.

Traditional approaches to scaling AI often rely on cloud-based processing power, but this model may be shifting. Public cloud providers have built their business models around offering scalable computing, massive storage and centralized data processing-services agentic AI systems need more of.

3. Interoperability Challenges

Multi-agent systems frequently need to interact with diverse data sources, APIs and other systems. Technologies like Model Context Protocol (MCP) servers demonstrate approaches to building tools that connect agents to databases and other services, but implementing these connections consistently presents technical challenges.



4. Data Privacy and Security Concerns

Agentic AI systems often require access to sensitive data, raising significant security and compliance concerns. Maintaining appropriate security controls becomes increasingly complex as agents interact with databases, APIs and user information. Like any digital asset in an organization, knowing their vulnerabilities or attack vectors is crucial to understanding how to protect them. The standard awareness document, "LLM OWASP Top 10," gives a guideline on what to look out for and what to do to protect your AI systems.

Analysis of threats to AI systems revealed that threat actors are experimenting with Gemini to enable their operations and using Gemini for research, troubleshooting code, and creating and localizing content.



Benefits of Agentic Al as a Strategic Technology Investment



TRANSFORMING CUSTOMER EXPERIENCES

In customer service environments, agentic AI has already achieved significant operational scale, orchestrating intelligence and automation across multiple service activities simultaneously. These systems demonstrate remarkable capability to analyze customer sentiment, review historical interaction data, access relevant company policies and respond to needs based on this comprehensive understanding. The result is a more personalized, effective service experience that addresses customer needs holistically rather than through fragmented interactions or siloed information access.

ENHANCING OPERATIONAL EFFICIENCY

Agentic AI delivers significant operational efficiency improvements across diverse sectors by automating complex, time-intensive tasks that previously required substantial human involvement. Autonomous AI agents can process and analyze massive data volumes faster than human operators, accelerating decision-making processes and streamlining workflows. In manufacturing environments, for example, agentic AI can autonomously manage and optimize production schedules in real time, minimizing downtime and maximizing output without requiring constant human supervision or intervention.





DRIVING INNOVATION AND DISCOVERY

Some of the most transformative agentic AI use cases emerge in domains requiring complex problem-solving, creativity and discovery processes. Scientific research and materials development provide compelling examples, where agentic AI capabilities extend far beyond identifying potential compound formulations. These systems can also evaluate supplier options based on priorities like cost efficiency or delivery timing, place orders for necessary materials, and coordinate the entire discovery-to-implementation process autonomously. This comprehensive approach accelerates innovation cycles while reducing the administrative burden on research teams.





TGS Tips

BEST PRACTICES FOR IMPLEMENTING AGENTIC AI

Successfully implementing agentic AI requires thoughtful planning, appropriate resources and effective ongoing management practices. Organizations can substantially increase their implementation success rates by following several key best practices throughout their AI adoption journey, addressing technical and organizational dimensions of effective deployment.

Set SMART goals: Align your agentic Al implementation with broader business priorities using the SMART framework: Specific (define what the Al system should accomplish), Measurable (establish quantifiable metrics for evaluation), Achievable (set realistic expectations given available resources and technology constraints), Relevant (ensure alignment with organizational needs), and Time-bound (create clear

☆☆☆ 288

Ø

Build an expert team: Assemble cross-functional teams (technical specialists, domain experts and end users), and engage change management experts / an AI council and executive sponsorship for a holistic approach. Incorporating these practical, human and financial aspects aids in successful implementation.

timelines for implementation phases and outcomes assessment).

Scaffold the decision space: Manage risks through continuous monitoring and well-defined boundaries around the AI system. Simultaneously, take a progressive approach to implementing autonomous systems while building decision-making capabilities and organizational confidence over time. This balanced approach to autonomy maximizes benefits while maintaining appropriate safeguards and human involvement.

Prioritize data quality and governance: Create clear policies for data collection, storage, usage and security to ensure compliance with relevant regulations and ethical standards while supporting the AI system's effectiveness. These governance frameworks should address technical aspects of data management and broader questions about appropriate data usage and protection.



_ _ _ _ _ _ _ _ _ _ _ _ .





Aireon, a satellite network operator who tracks air traffic control data, is an early adopter of agentic AI for security operations (SecOps). They sell software to air navigation service providers and other companies in the aerospace industry. Aireon's 66-satellite network spans the globe but operates with fewer than 200 employees, according to the company's CISO, Peter Clay. With a corporate computing and cybersecurity staff of just five people, threat hunting across 130 million to 150 million log events daily is no small task.

Compounding this challenge is a general SecOps skills gap in the industry. New SecOps professionals enter the industry annually, but mid-level talent with five to 15 years is scarcer. Starting in 2024, Clay addressed these gaps with an open source LLM, WhiteRabbitNeo, to identify and remediate security vulnerabilities. Aireon also employed a product from WhiteRabbitNeo commercial sponsor Kindo that provided a framework to automate DevSecOps workflows using runbooks and AI agents. Rather than needing a threat analyst (five to seven years of training) to hunt and deliver valuable data, the LLM finished in just a couple of hours. Since deploying WhiteRabbitNeo and Kindo, Clay's team has caught two active advanced persistent threats and shut them down before the attacks could advance.

An Aireon engineer built a set of AI agents with Kindo to gather and query the company's security plan documentation for developers as they work. The SecOps team uses this same set of agents to upload new versions of the documentation, track queries and identify which developers need help implementing security controls-freeing up time and providing focused feedback.

To measure the effectiveness of WhiteRabbitNeo and Kindo, Aireon uses another agentic Al tool from Derive for finance automation, risk management and compliance. Derive estimates that the Al agents have helped the company avoid \$2 million in risk exposure.³

All information shared herein was accessed from public sources as indicated and is not indicative that the named entity is a TEKsystems Global Services client nor that the work was performed by TEKsystems Global Services.





TGS PARTNER PORTFOLIO

- · As an AWS Premier Tier Services Partner, TEKsystems covers the full spectrum of Amazon Web Services (AWS) initiatives. From design, migration and implementation to adoption and improvement, continuous integration and delivery (CI/CD) to infrastructure as code, Lean-Agile and more-we're there.
- As a Google Cloud Premier Partner, we support the full spectrum of delivering Google Cloud initiatives, from design, migration and implementation to adoption and improvement, covering CI/CD, infrastructure as code, Lean-Agile, data analytics, AI, ML and Gen AI.
- As a Microsoft Solutions Partner, we bring gualified expertise and deep experience to help you maximize ROI and achieve real value. From discovery and design to adoption and improvement-we'll tailor our solutions to meet your needs and help you stay ahead of what's next.
- As a **Red Hat Premier Business Partner**, we provide qualified technical leadership, open source expertise and scale to help you get the most out of your Red Hat products-no matter where you are in your modernization journey.
- With 30-plus SnowPro certified architects, our Snowflake Elite Partner status highlights our proven skills and experience to help you leverage Snowflake's innovative technology and achieve datadriven results.
- As a **ServiceNow Elite Partner**, we bring experience and subject matter know-how to help you drive your ServiceNow initiatives. From implementation to optimization-we'll tailor our services to help you stay ahead of the curve and accelerate the adoption of ServiceNow solutions.
- · As a Salesforce Summit Partner, our experienced team can help you maximize the value of the platform to harness each opportunity. With technical insight, delivery excellence and the deepest bench in the game (500-plus Salesforce-certified pros), we can help you transform your sales, service and communities into true power players.

IN GOOD COMPANY

Transformational technologies demand equally transformative partnerships. We offer full-stack capabilities coupled with depth and diversity of experience in leading platforms that help organizations grow, innovate and thrive.

The views and opinions expressed in this publication are those of the authors and do not necessarily reflect the views of TEKsystems Global Services, LLC or its related entities.





Alvaro Albanez

Alvaro Albanez is an expert in cybersecurity, cloud migrations and DevOps with 20 years in the IT industry. He engages with TEKsystems Global Services customers in healthcare, finance and retail, navigating complex cloud environments to create secure workloads and applications with high-performing teams.



Armando Franco

Armando is a seasoned technology leader. He thrives at the intersection of strategic leadership and the delivery of complex multicloud solutions. With over 22 years of experience in the IT industry, he has dedicated his career to driving innovation and excellence in technology. Throughout his career, Armando has been consistently driven by a passion for transforming businesses through technology. For the past six years, he has used his expertise to elevate TEKsystems Global Services.



Ramesh Koovelimadhom

Ramesh Koovelimadhom, Director of Google Cloud Partner Solutions, **TEKsystems Global Services**

With over 25 years of experience in the industry, Ramesh is responsible for driving growth with Google technologies. Aligned closely with our practice and sales leadership, he interfaces with multiple layers of our client's technology and business management to identify, position and deliver business outcomes. He has successfully led several digital transformation engagements and helped clients in bridging the strategy-execution gap and resetting the culture in the IT organization, translating the strategy to everyday plans and reorganizing costs to grow stronger.

Associate Director of Cloud Security, TEKsystems Global Services

Director of Business Modernization Services, TEKsystems Global Services





About TEKsystems[®] and **TEKsystems Global Services**

We're a leading provider of business and technology services. We accelerate business transformation for our customers. Our expertise in strategy, design, execution and operations unlocks business value through a range of solutions. We're a team of 80,000 strong, working with over 6,000 customers including 80% of the Fortune 500—across North America, Europe and Asia, who partner with us for our scale, full-stack capabilities and speed. We're strategic thinkers, hands-on collaborators, helping customers capitalize on change and master the momentum of technology. We're building tomorrow by delivering business outcomes and making positive impacts in our global communities. TEKsystems and TEKsystems Global Services are Allegis Group companies. Learn more at TEKsystems.com.



Don't miss Alvaro Albanez, Armando Franco and Ramesh Koovelimadhom on The Agile World podcast. In a three-part series, host, author and business expert Greg Kihlström sits down with Alvaro, Armando and Ramesh to discuss how agentic AI uses autonomous problem-solving and large language models to boost innovation and business success.



Be in the Know

Check out previous issues and know what's next: Version Next, Now.



Sources

1. 2025: Agentic And Physical Al—A Multitrillion-Dollar Economy Emerges, Forbes 2. The Fascinating History of Agentic Artificial Intelligence, Medium

3. <u>CISO Enlists Agentic Al Reinforcements for SecOps</u>, TechTarget



TEKSYSTEMS.COM I AGENTIC AI ISSUE I SPRING 2025

