The Agent Advantage

Navigating Enterprise AI's New Strategic Imperative

Mark Beccue | Principal Analyst, AI
ENTERPRISE STRATEGY GROUP

January 2026

This eBook from Enterprise Strategy Group (now Omdia) was commissioned by Cognizant and Google Cloud and is distributed under license from TechTarget, Inc.

Introduction

The enterprise technology landscape is witnessing a seismic shift, with AI agents emerging as the new frontier of enterprise AI. Far from being just another technological trend, AI agents have rapidly ascended to become strategic imperatives for organizations worldwide.

As organizations increasingly recognize the transformative potential of AI agents—from autonomous decision-making systems to sophisticated multi-agent collaborations that tackle complex problems—they're setting ambitious expectations for operational efficiency and workflow automation. Yet this enthusiasm is tempered by clear-eyed recognition of challenges, including limited internal expertise, deciding between pre-built agents and those made from scratch, compliance concerns, integration complexities, orchestration, monitoring, and observability. This eBook explores the current state of the AI agent market, examining both the tremendous opportunities and significant obstacles organizations face as they navigate this rapidly evolving landscape, while highlighting how, together, Cognizant and Google Cloud are the ideal partners to position organizations for success in this new era of agentic AI.

KEY FINDINGS/CONTENTS

Back to Contents

State of the Market: AI Agents

AI Agents Become the Priority for Enterprise AI

There is a titanic shift in enterprise AI strategy to AI agents. Recent research from Enterprise Strategy Group (now Omdia) showed that 80% of organizations said AI agents are their top or one of their highest priorities compared to other AI initiatives.

Reflecting this priority is organizations' investment in the technology: Budget allocation for AI agents is outpacing historical generative AI spending patterns. AI agent budgets are ramping up more quickly than generative AI budgets did in comparable stages of the market. Forty-two percent of respondent organizations are allocating $1 million or more to AI agents over the next 12 months, whereas generative AI, at a similar stage (one year after first commercial vendor offerings were launched), saw only 12% of respondent organizations with budgets of $1 million or more. AI agent initiatives are important enough to organizations that 72% of AI agent budgets are either new standalone investments or being redirected from other non-AI related initiatives.

AI Agent Momentum Shuffles Technology Investments

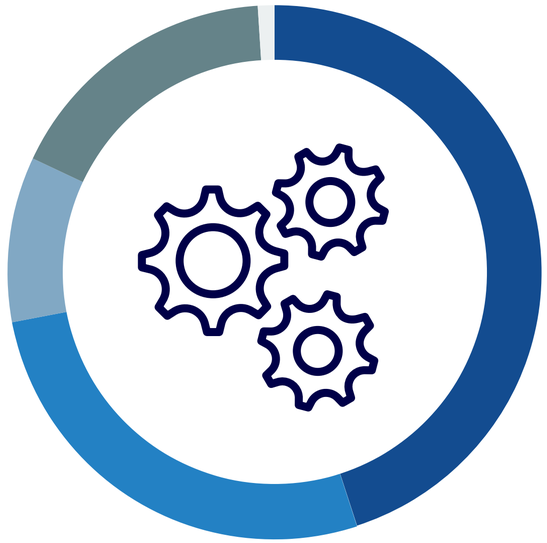

- **45%** Using a completely new standalone budget
- **27%** Reallocating funds from other technology or innovation budgets
- **10%** Reducing existing budgets in other AI areas (e.g., taking from GenAI and traditional AI budgets to fund AI agents)
- **17%** Using budget from part of a larger AI or digital transformation initiative

AI Agent Budgets Surge

- **5%** Less than $100,000
- **24%** $100,000 to $499,999
- **27%** $500,000 to $999,999
- **29%** $1 million to $4.999 million
- **13%** $5 million or more

Source: Enterprise Strategy Group (now Omdia) Research Report, AI Agents: The Game-changing Generative AI Use Case, August 2025.
All Enterprise Strategy Group research references and charts in this eBook are from this report unless otherwise noted.

Back to Contents

Expectation for AI Agents Is High

While generative AI's wide-ranging potential elicited plenty of unrealistic expectations, AI agents are a much more pragmatic interpretation of generative AI's capabilities, as the market is primarily focused on deploying AI agents to create and accelerate business efficiencies in productivity, automating business workflows and processes, and improving and speeding decision-making. Forty-six percent of organizations strongly agreed AI agents will improve their organization's efficiency and automation. Additionally, 45% strongly agreed AI agents can streamline complex tasks, significantly reducing manual workloads.

Organizations Have High Expectations for the Business Impact of AI Agents

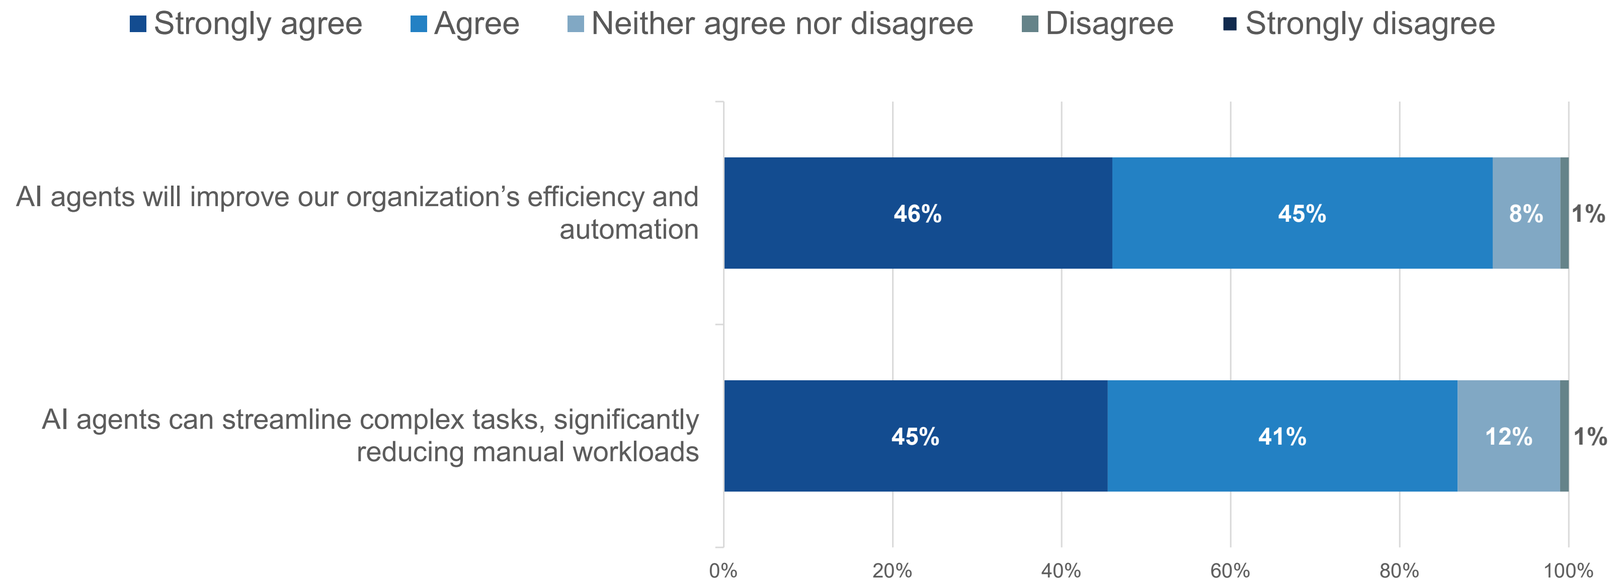

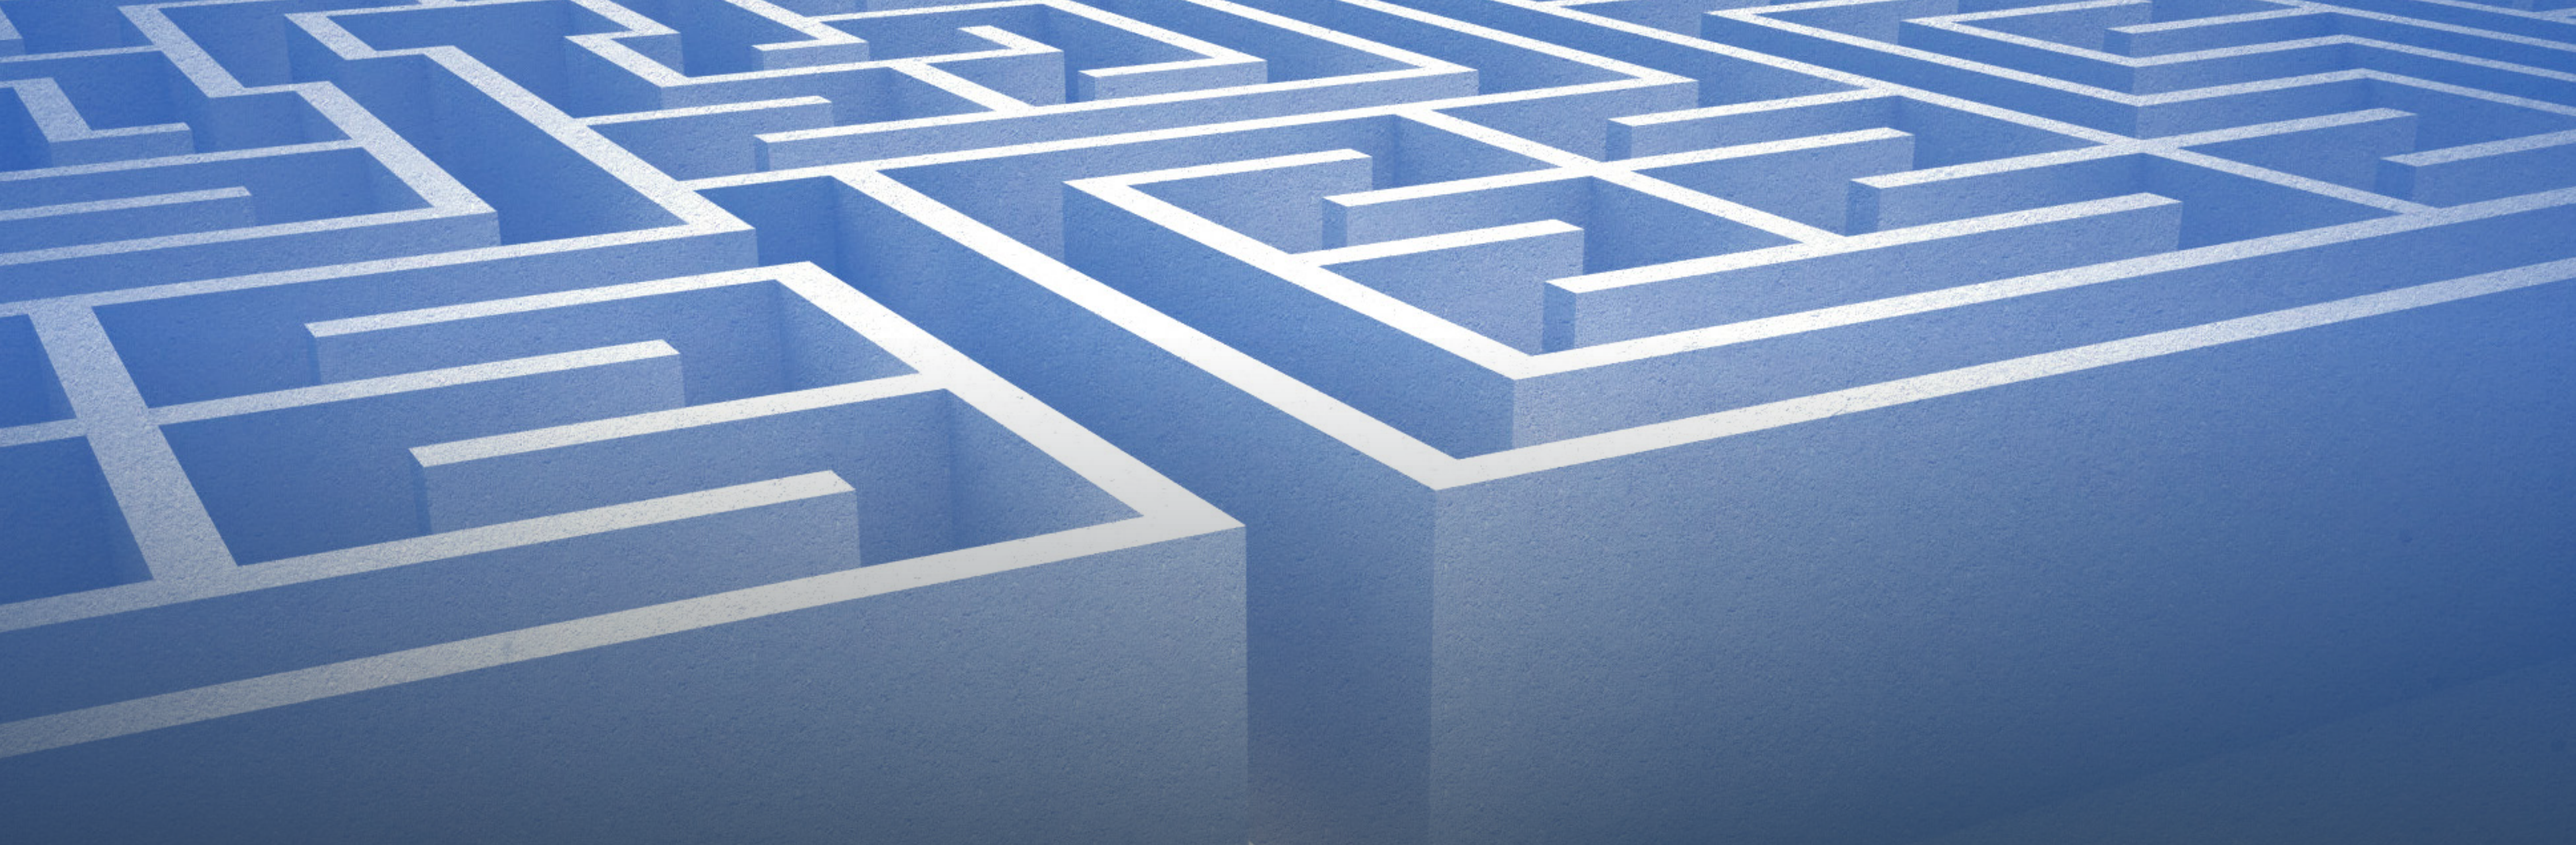

Challenges for AI Agents

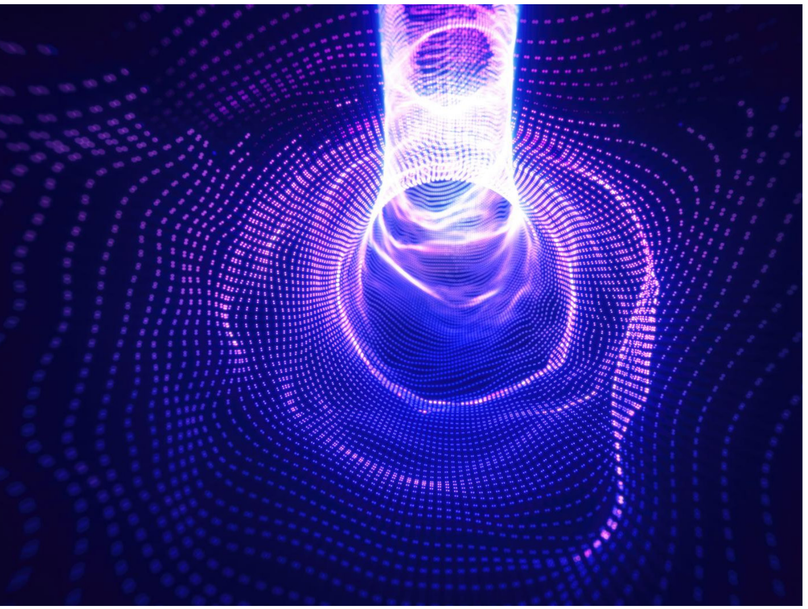

Agent Bottlenecks: Culture, Complexity, and Acting on Practical Use Cases

Despite the enthusiasm for AI agents, organizations also understand there are challenges to building and managing them, not the least of which are cultural issues primarily centered around employee resistance due to job security fears, skills gaps in AI collaboration, and trust issues with automated decision-making systems. These challenges are compounded by disrupted communication patterns, leadership gaps in managing human-AI teams, and ethical concerns about bias and accountability in AI.

Building sophisticated agent systems is a complex undertaking. According to Enterprise Strategy Group (now Omdia) research, 81% of organizations agreed that scaling AI agents into full production is more complex than they initially expected.

Many organizations have struggled with latching onto practical use cases, despite being clear about their business drivers for AI agents. When asked what their most important drivers for the use of AI agents were, 39% of respondents said to increase productivity, 38% said to improve and/or automate processes and workflows, and 31% cited the ability to enhance customer experience and engagement.

Top 5 Most Important Business Drivers for the Use of AI Agents

39%

Increase productivity

38%

Improve and/or automate processes and workflows

33%

Improve decision-making speed and accuracy

31%

Enhance customer experience and engagement

26%

Strengthen security, risk management, and compliance

Back to Contents

Security and Compliance Concerns

Organizations relish the opportunities to leverage agents, but they are wary of what risks such autonomous applications pose. In Enterprise Strategy Group's (now Omdia) research, respondents said security and compliance concerns are the most significant challenge they face in implementing AI agents. When asked which risks related to AI agents cause the most concern, the top three answers focused on security and compliance: Part of their security and compliance concerns is how to properly develop monitoring and observability. Hard questions leaders must grapple with include:

- 37% data privacy compliance
- 34% security vulnerabilities such as exposure to cyberattacks or breaches due to AI vulnerabilities
- 29% failure to meet industry regulation or government standards

Part of their security and compliance concerns is how to properly develop monitoring and observability. Hard questions leaders must grapple with include:

- How do we enforce compliance in environments where agents are empowered to commit, test, and deploy code autonomously?
- How do we ensure generated outputs respect architectural principles, security policies, and industry regulations?
- How do we prevent documentation, decisions, and artifacts from drifting apart, causing confusion and risk?
- How do we maintain systems that have deterministic outputs arrived at with probabilistic reasoning?

Top 3 Concerning AI Agent-related Risks

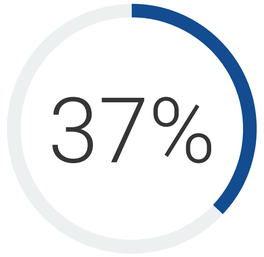

Data privacy (i.e., risks related to handling sensitive data in compliance with privacy laws)

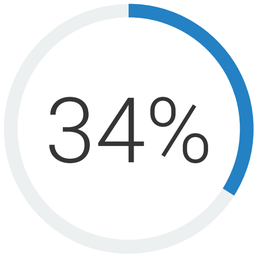

Security vulnerabilities (i.e., exposure to cyberattacks or breaches due to AI vulnerabilities)

Compliance and regulatory risks (e.g., failure to meet industry regulations or government standards)

Back to Contents

System Integration and Deployment Complexity

Agents tap a range of data sets and software applications and consequently introduce complexity. AI agents will test organizations due to the difficulty in creating interoperability between disparate software and systems. These challenges will become more prominent as organizations seek to leverage AI agents that access and communicate with other AI agents outside of the organization's direct control. Agents are capable of reasoning and learning across tasks and environments with significant autonomy. Multi-agent systems are designed to collaborate, negotiate, and solve problems. Agents can be designed to continuously learn from their experience and then make real-time decisions in dynamic environments. In short, they are complex systems to build and operate.

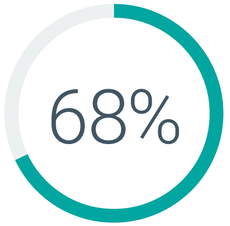

of organizations agreed they face significant challenges integrating AI agents into existing systems and processes.

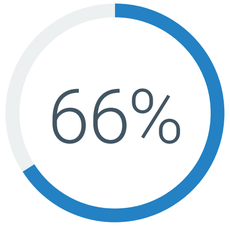

agreed bridging disparate systems and data sources to enable AI agents is a significant challenge.

AI Agents Will Pose Deployment Challenges

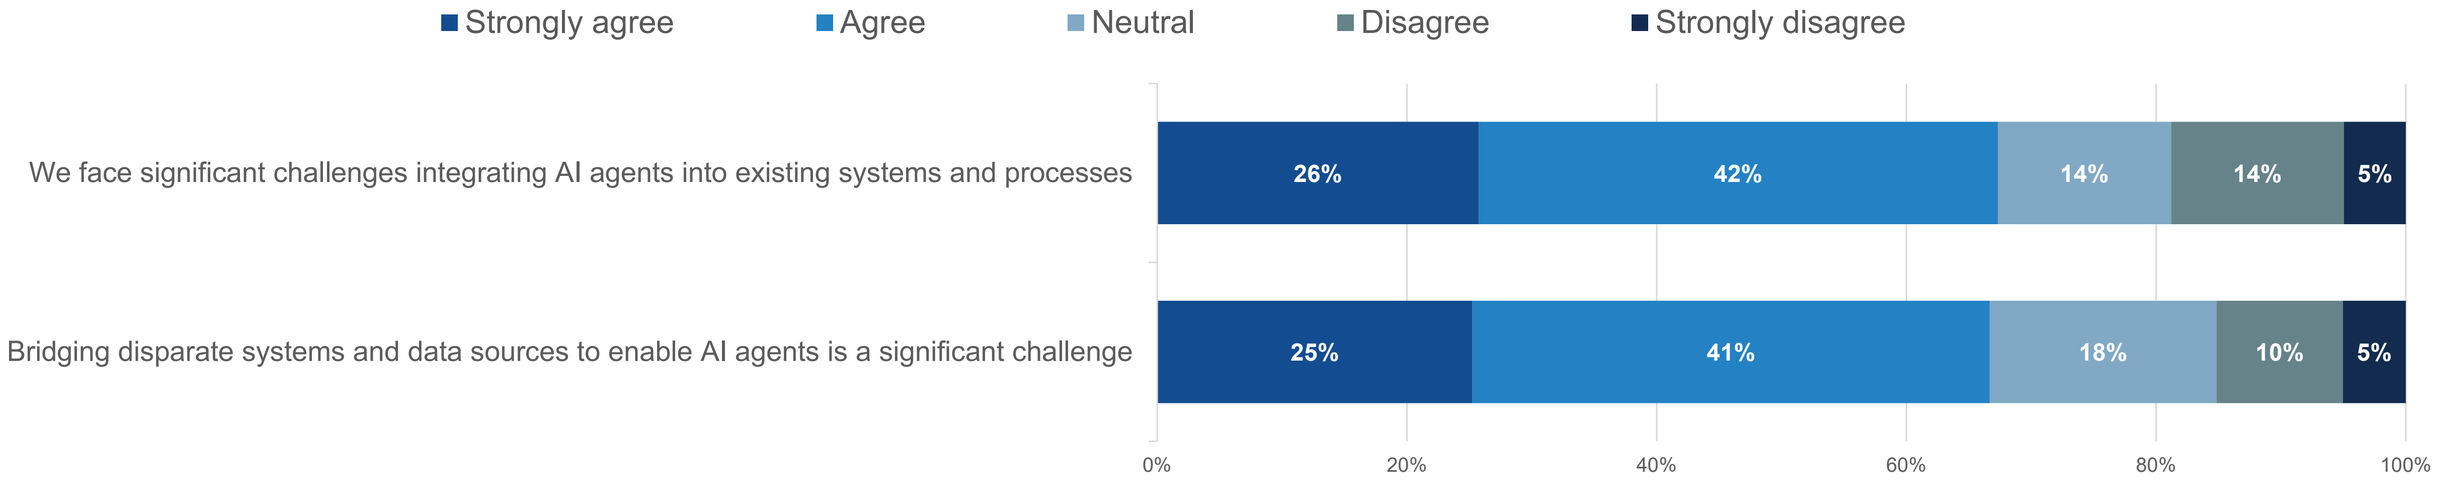

Back to Contents

The AI Agent Approach: Best Practices

Create AI Centers of Excellence

AI Centers of Excellence (AI CoEs) are emerging as the essential backbone for sustainable, enterprise-grade adoption of agentic AI. AI CoEs will define the rules of engagement for the agentic era and orchestrate expertise across technology, business, compliance, and human factors, with executive leadership at the core given the business-wide impact. Their remit extends beyond traditional infrastructure to encompass models, connectors, context frameworks, and trust controls as first-class assets, fundamentally reshaping how organizations approach AI integration.

Rethinking Top-down Approach: A top-down, bottom-up approach incorporates excecutive-driven direction, vision, and leadership, with AI solution advocates and pioneers who are experimenting and innovating at the use case/tactical level.

New Roles: C-suite owners as executive drivers, AI policy specialists for ethics and compliance, model lifecycle owners for enterprise model management, context engineers, trust and observability leaders who are ensuring explainability across workflows, and change managers bridging culture and skills uplift.

These shifts ensure teams work confidently alongside AI rather than being merely governed by it. Organizations should operate their AI CoE like a small product company, establishing opinionated defaults, encouraging open contribution, maintaining standards in code, and embracing federated ownership.

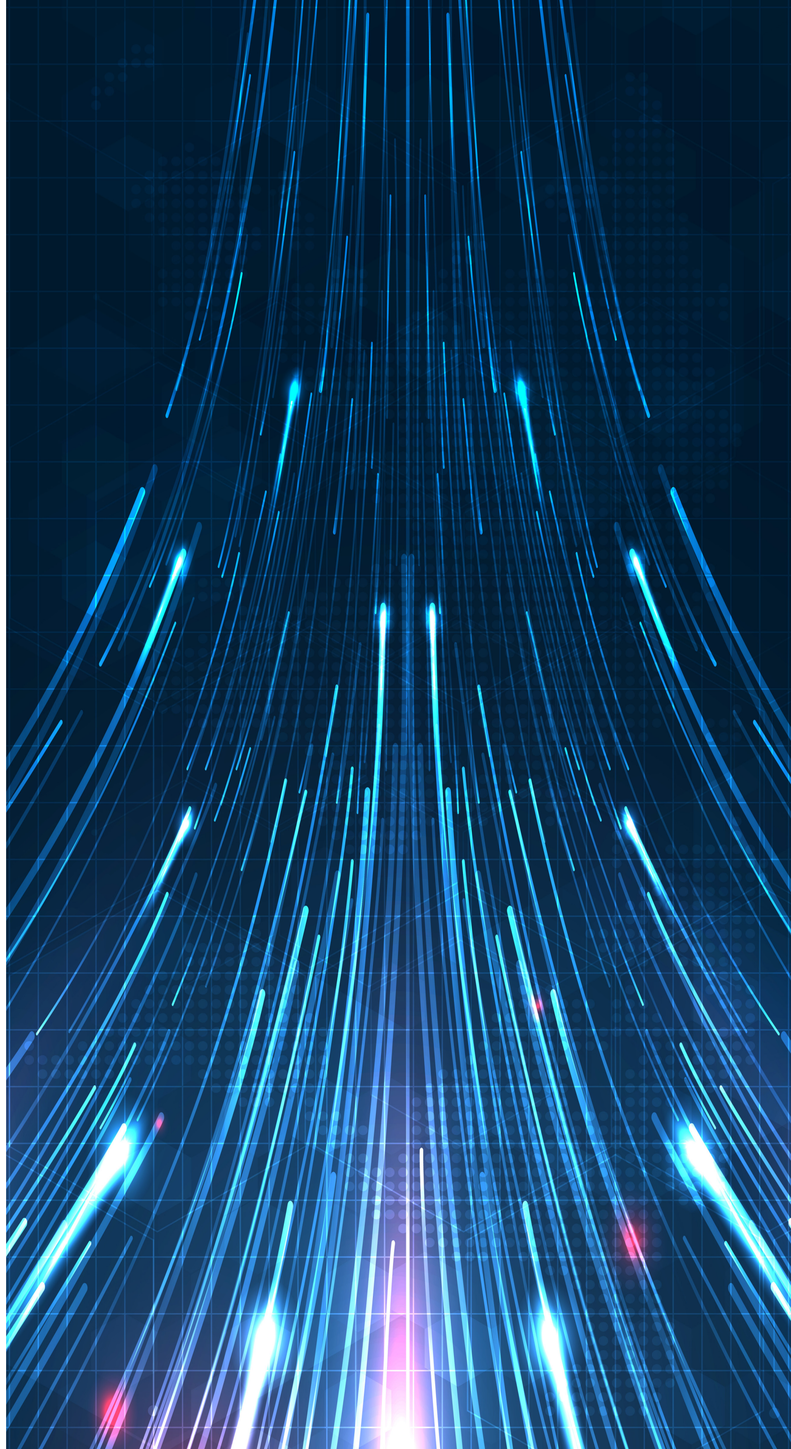

Identify and Prioritize High-value Use Cases

Innovation with agentic AI depends on a classic business discipline principle that takes the form of the question: What problems are we trying to solve? Organizations should start by evaluating business processes for viability of agentic automation.

Key Principles for High-value Use Cases

Context and Governance: Embed organizational knowledge and governance into AI systems to ensure agents operate within defined standards and regulations, moving beyond pilots to enterprise-scale impact.

Ethical AI: Prioritize transparency, fairness, and regulatory compliance in the design and deployment of agentic AI solutions.

Business-driven Outcomes: Focus on delivering proven, real-world functions that reduce manual labor, enhance operational excellence, and achieve tangible business value for clients.

Back to Contents

Rapid Agent Development: From Proof of Concept to Proof of Value in Days

A primary obstacle to enterprise AI adoption is the inertia of traditional development cycles. Lengthy proofs of concept, which can take months to complete, often delay tangible value and cause promising initiatives to lose momentum. To overcome this, Cognizant and Google Cloud have pioneered a new engagement model designed for velocity: **Rapid Agent Development,** powered by Google's Gemini suite of products and capabilities.

This methodology reframes the goal from creating a "proof of concept" to delivering a "proof of value"—a production-ready agentic AI solution—in a compressed time frame of just a few days.

A Framework for Selecting the Ideal Pilot

The success of an **AI Sprint** for **Rapid Agent Development** hinges on a disciplined approach to use case selection. While the build process is accelerated, the initial scoping is highly strategic. The framework focuses on identifying opportunities that are:

- **High in business value:** The solution must address a clear and measurable business pain point, ensuring the outcome is immediately relevant to stakeholders.
- **Adjacent to core business functions:** The ideal pilot use case is one that provides significant operational improvement without directly impacting mission-critical systems, thereby mitigating risk for the initial deployment.
- **Broad in user impact:** The target solution should have a significant user base (e.g., 500+ employees) who would use it regularly, guaranteeing that a successful pilot will generate immediate and widespread adoption data.

The Outcome: A Production-ready Agentic AI Solution Package

Unlike a traditional hackathon that produces a prototype, the AI Sprint for Rapid Agent Development delivers a complete, deployable package. At the conclusion of the sprint, the customer receives not just a working agent deployed on a production environment but also a suite of Gemini-generated assets, including user documentation, a deployment guide, and an executive-level demonstration. This model serves as a powerful catalyst, de-risking the AI investment, generating crucial organizational momentum, and providing a scalable, repeatable blueprint for driving the agentic enterprise forward.

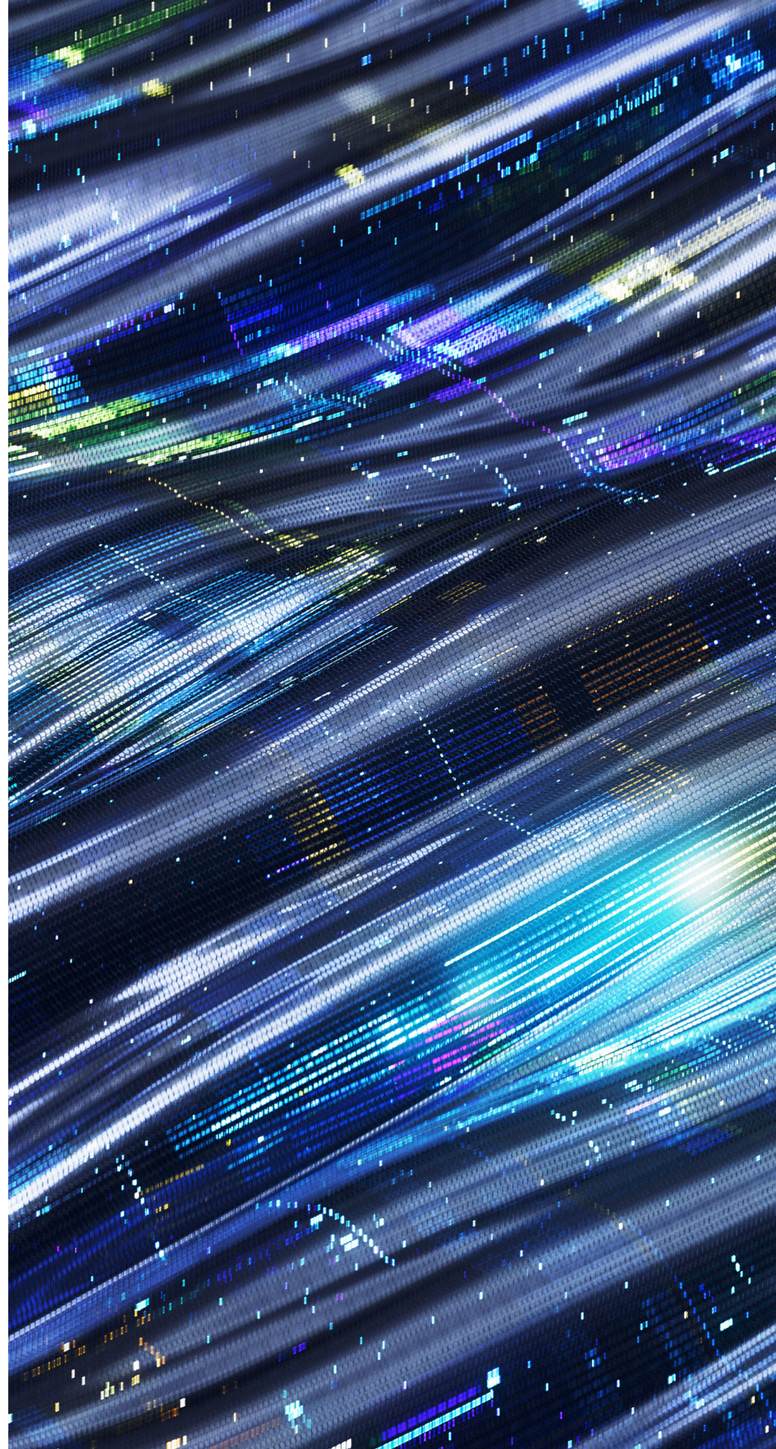

Develop a Culture of Change: People and New Skills

For AI agents to succeed, people and new skills are a critical component. Organizations should promote a culture of learning and innovation while building safeguards and support so teams aren't afraid to take risks.

The rapid pace of AI change means job titles and roles in many organizations will eventually be more fluid, akin to how roles within startups are structured. This mindset—shifting individuals' working attitudes to align with the needs of the organization, is critical. From a top-down perspective, leaders need to focus on clear communication, clear direction and vision, and be willing to be flexible and adapt the strategy when facing more change. In terms of bottom-up strategy, focus on identifying individuals who can lead and galvanize agentic AI adoption and change and advocates who showcase the art of the possible.

AI agents have intensified the focus on human-AI collaboration. When asked in which areas their organization has AI agent skills gaps, 33% of respondents said human-AI collaboration, the top response given. For upskilling and reskilling, organizations need to think about training resources to be the critical human-in-the-loop guardrail. Perhaps just as importantly organizations need to develop context engineering.

The New Critical Role: Operationalizing Context Engineering

Success in the agentic era requires a new architectural layer: Context. With 33% of organizations citing "human-AI collaboration" as a top skills gap, Context Engineering emerges as the essential discipline to bridge enterprise intent with agent behavior. Unlike prompt engineering, which is tactical, Context Engineering is strategic. It is the practice of designing the environment so agents "represent enterprise knowledge, align with human intent, and reflect company work-culture."

To operationalize this, the AI Center of Excellence must mandate three core responsibilities for the Context Engineer:

1. Knowledge Grounding: Anchoring agents in "core business systems" to ensure accurate, non-hallucinated responses.
2. Cultural Alignment: Codifying "human goals" and "company work-culture" into system instructions, ensuring agents negotiate and collaborate rather than just process data.
3. Governance Engineering: Answering the "hard questions" of security by embedding "opinionated defaults" and compliance guardrails directly into the agent's logic.

Feature	Prompt Engineering	Context Engineering
Scope	Tactical (Task-based)	Strategic (System-based)
Focus	"Write a good email."	"Who are we, and what are our policies?"
Outcome	A single valid response.	Trust, alignment, and reduced risk at scale.

Cognizant is deploying 1,000 Context engineers, powered by ContextFabric™, to industrialize agentic AI. Learn more about Cognizant's strategic investment in context engineering.

Governance

Frameworks are needed to manage AI agents responsibly. This includes governance processes for data provenance, model explainability, decision-making transparency, as well as establishing an ethical framework for autonomous agents.

Back to Contents

Best Approach: Find an Expert Partner

Strategic partnerships are not just beneficial but essential for successful enterprise implementation of AI agents.

- 65% of organizations that have started agent initiatives already partner with AI vendors or service providers to deploy AI agents.
- 91% of organizations reported that finding the right external partner has been critical to their AI agent deployment success. This overwhelming consensus underscores a fundamental truth: AI agents represent a paradigm shift that requires specialized expertise, proven methodologies, and deep technical knowledge that most organizations simply don't possess in-house.

The ideal AI agent partner brings a comprehensive suite of capabilities that extends far beyond basic technical implementation. They:

- Posess hands-on experience operationalizing multi-agent systems at enterprise scale.
- Demonstrate the strategic acumen to assess an organization's AI maturity and craft tailored approaches accordingly.
- Understand the full AI agent lifecycle—from security and compliance frameworks to orchestration, monitoring, observability, and DevSecOps integration.
- Address the human elements of transformation, helping organizations navigate the people, processes, and cultural changes that AI agents inevitably bring.
- Offer flexibility in their approach, providing both pre-built AI agent solutions for rapid deployment and the platforms and tools necessary for custom development.
- Become an extension of the organization's capabilities, accelerating time to value while mitigating implementation costs by combining state-of-the-art, enterprise-grade infrastructure with services that enable quick prototyping, customization and scaling.

Back to Contents

Why Cognizant + Google Cloud

Cognizant's Specialized Expertise

In partnership with Google Cloud, Cognizant delivers not just AI solutions but the operating model transformation required to turn pilots into competitive advantage. Cognizant and Google Cloud have established a comprehensive strategic alliance specifically focused on accelerating AI adoption for enterprises. This partnership combines Cognizant's deep industry expertise and consulting capabilities with Google Cloud's cutting-edge AI infrastructure and tools, creating a powerful foundation for AI agent deployments.

Cognizant advantages

- Experience and expertise in operationalizing AI agents at scale. Experience in guiding people, processes, and culture changes AI agents bring.
  - Experience creating AI CoEs
  - Leverage Cognizant's 1,000 Context Engineers
- Industry-specific experience to bring the context of drivers and barriers to bear in developing agent systems tailored to your organization's specific needs. Cognizant brings deep vertical expertise across industries such as financial services, healthcare, manufacturing, and retail, enabling tailored AI agent solutions that address specific business challenges and regulatory requirements.
- Understanding of the complete lifecycle of AI agents and the components that need to be addressed, including security and compliance frameworks, orchestration, monitoring, observability, and systems integration. From strategy and design to deployment and ongoing management, Cognizant provides comprehensive services that cover the entire AI agent lifecycle, including governance, security, compliance, and change management.
  - **Cognizant Agent Foundry.** Specifically designed to help enterprises design, deploy, and orchestrate autonomous AI agents at scale, providing reusable assets and implementation services. Cognizant develops high-value agentic use cases by focusing on cross-functional applicability and enterprise-grade implementation through a four-step process: **Ideate, Design, Engineer, and Scale.** This approach emphasizes prioritizing the right use cases, engineering the fit-for-purpose solution, integrating with enterprise systems using reusable components, and, finally, scaling deployments across geographies with governance, observability, and ROI tracking to ensure business value and mitigate risks.

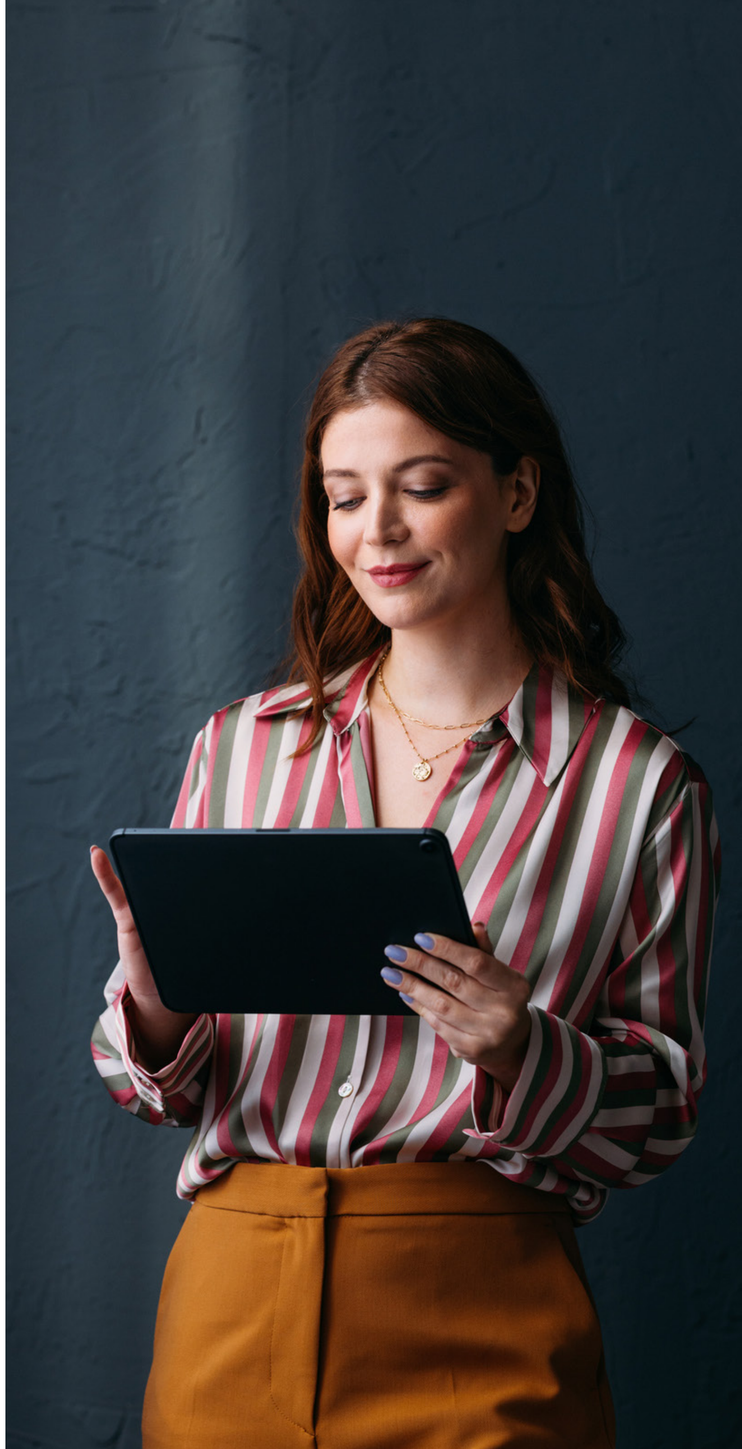

Google Cloud

Google Cloud provides a secure, scalable, full-stack platform to build, deploy, and manage sophisticated AI agents.

- **Advanced AI Models:** At the core, Gemini models are multimodal, meaning they can natively reason across text, images, and data. This enables agents to understand complex, real-world information, not just text commands.
- **Professional AI Platform:** Vertex AI Agent Builder is the "professional workshop" for the entire AI lifecycle. It provides the tools to build, deploy, and govern agents, with robust grounding capabilities to securely connect to your enterprise data for accurate, relevant responses.
- **AI for the Workforce:** Gemini Enterprise serves as the secure "front door" for all employees. It integrates AI directly into the flow of work by securely connecting to your core business systems, including Google Workspace, Microsoft 365, and Salesforce.
- **Global Infrastructure:** This all runs on Google's planet-scale, secure-by-design infrastructure, optimized with custom AI accelerators for unparalleled performance and efficiency.

Cognizant + Google Cloud

The Cognizant Google Cloud Business Group is a powerful alliance built on top executive commitment and proven results, as evidenced by numerous industry awards and recognition. This dedicated business group focuses exclusively on Google technologies and go-to-market strategies, creating a specialized ecosystem that accelerates cloud transformation for clients across industries. The partnership's strength lies in its comprehensive approach—blending culture, talent, and technology to foster innovation while leveraging Cognizant's deep experience in data, AI, and modernization alongside Google Cloud's robust data expertise. This unique combination enables the creation of industry-specific strategies that not only drive immediate business results but also establish foundations for continuous innovation on Google Cloud's secure infrastructure.

The tangible power of this partnership is demonstrated through the development of custom Google-architected agents which optimize client experiences and deliver measurable outcomes. Cognizant's track record of executing large transformational projects—backed by ongoing innovation investments and systematic upskilling of digital talent—provides clients with the confidence that complex cloud migrations and modernization initiatives will be successful. This proven experience, combined with Google Cloud's cutting-edge infrastructure and data capabilities, creates a synergistic relationship that enhances client relevancy in their respective markets, improves customer experiences, and delivers tangible business benefits that extend far beyond traditional cloud adoption. Learn more about the Cognizant and Google Cloud strategic partnership.

Back to Contents

Conclusion

Future: The Agentic Enterprise

The future of AI is moving toward the agentic enterprise as a fundamental transformation where AI agents become, in Cognizant's view, "embedded, trusted participants in how work gets done and transformation is managed." This represents a paradigm shift from AI as a tool to AI as a collaborative partner working alongside humans in continuous, adaptive operations. The future enterprise will feature autonomous, goal-oriented systems that operate independently while remaining aligned with business objectives, enabling real-time decision-making and instant responses to market changes, customer needs, and operational challenges.

Central to this vision is a human-AI collaboration model that doesn't replace human workers but creates new forms of partnership where humans provide creativity, strategic thinking, and ethical oversight, while AI agents handle routine tasks, data analysis, and process optimization. Through their partnership, Cognizant and Google Cloud see enterprises building agentic capabilities using industrialized agent templates, domain-specific models, and orchestration systems that manage multiple AI agents working toward common goals. Leveraging deep business domain expertise across industries, these AI and agentic systems are now fundamentally transforming business processes, rapidly reducing process time and operational costs while delivering unprecedented efficiency gains. This composable, platform-agnostic architecture enables continuous agent-driven transformation, moving beyond periodic initiatives to ongoing evolution where AI agents continuously optimize processes, identify opportunities, and implement improvements. By combining industry-specific knowledge with advanced agentic capabilities, enterprises can deploy AI agents that understand the nuances of their particular business domains—whether in financial services, healthcare, manufacturing, or retail—enabling them to make more intelligent decisions and drive more impactful process improvements. This approach ultimately creates more responsive, efficient, and innovative enterprises that can adapt and evolve at the speed of business.

Back to Contents

ABOUT COGNIZANT

Cognizant (Nasdaq-100: CTSH) engineers modern businesses. It helps its clients modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. Together, Cognizant is improving everyday life. See how at www.cognizant.com or @cognizant.

LEARN MORE

Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.