BARC

Innovating with Agentic AI

How enterprises are reshaping strategy with AI-powered agents

Author: Shawn Rogers

Publication: July 2025

Research sponsored by: Endava PLC

Table of Contents

Executive Summary

Agentic AI shifts artificial intelligence from insight to autonomous action. As adoption accelerates in 2025, you may soon manage hundreds – or thousands – of goal-driven agents automating tasks from customer service to clinical coding. This research note defines AI agents, the technology's key capabilities, market trends, and the strategy companies should follow prior to deployment.

This research note utilizes findings from several global research reports published by BARC. Each report has 250-500 qualified global respondents. This research shows 32% of firms already run agents in production and another 26% are piloting initiatives, yet only 21% possess the governance, security, and data foundations large-scale use demands. Inadequate controls raise the risk of biased outputs, privacy breaches, and operational threats, but companies that act early gain faster time-to-value and significant savings – such as the projected $35 million a top 10 pharmaceutical customer is capturing by automating clinical coding with AI agents.

BARC research shows 32 percent of firms already run agents in production.

Agentic systems blend machine learning, natural language processing, and small or domain language models to perceive context, decide, and execute without human prompting. Multi-agent architectures use reflection loops, so reviewer agents test and refine code or decisions, keeping every action auditable. Emerging standards like model context protocol (MCP) simplify tool integration, while cloud GPUs supply the horsepower for real-time responses at scale. Yet, autonomy magnifies risk; strong governance, security, and human-in-the-loop design remain non-negotiable.

Execution starts with establishing a solid foundation across people, process, and technology. To capture benefits and avoid missteps, executives should focus on these priorities:

- Appoint accountable AI leadership to own strategy, budget, and outcomes.
- Harden data quality, access, and lineage before scaling autonomous workflows.
- Embed governance agents for compliance, ethics, and risk monitoring in every multi-agent stack.
- Select small or domain language models to improve context and control costs.
- Define clear, measurable goals so each agent can self-evaluate and report progress.

With these foundations in place, you can scale from simple pilots to a resilient, enterprise-wide agent workforce that accelerates revenue and reduces operational costs.

An Introduction to Agentic AI

This innovative technology can provide a quick onramp to leveraging artificial intelligence (AI) for automating processes that support revenue growth across enterprise business landscapes. It's likely that most large companies will soon find they have hundreds, if not thousands, of AI agents operating autonomously throughout the organization.

Understanding Agentic AI

Let's start by defining agentic AI and agents to better understand why they are important. Agentic AI is a sophisticated form of AI; it's often defined by its autonomy and agency, meaning agents make decisions and work toward achieving specific goals and outcomes. Agentic AI can execute simple tasks or transactions, but can also work its way through complex decision trees to deliver highly valuable outputs and action. An agent's "perception" of its environment and the context surrounding it allows it to work in a human-like fashion to solve problems and execute tasks while operating autonomously at scale and with limited oversight. AI agents can become smarter over time by learning to optimize their efficiency. Agents are the next step beyond simple rule-based, policy-driven robotic process automation (RPA) systems to take insight-driven action.

Agentic Importance and Impact

Agentic systems are becoming critical components of successful AI strategies. Agents leverage proven technologies like machine learning (ML), natural language processing (NLP), and generative AI (GenAI), and can scale well beyond manual human-driven processes. The economic impact of AI has been estimated to add $15.7 trillion to the global economy by 2030[1]. While these market predictions vary widely, it's clear AI is having a significant impact economically and it would be a critical error for companies not to embrace this opportunity.

The economic impact of AI has been estimated to add 15.7 trillion dollars to the global economy by 2030.

Agent Capabilities

Because AI agents perceive, decide, and act independently, agent capabilities deliver a different type of value to the enterprise than traditional AI insights. Traditional AI insights often stop short of taking autonomous action leaving these steps to humans. The time gap created in many of these interactions reduces the value or loses opportunity altogether. Agents can share information with each other, learn from that information, and leverage it to take accurate action based on predictive models and data analysis.

Traditional AI versus Agentic AI

Traditional AI systems are more restricted in that they follow pre-programmed rules to accomplish specific tasks. Agentic AI learns from interactions, understands context, and can adapt its behavior. This functionality in agents is more sophisticated and can execute more meaningful tasks.

[1] PwC: "Global Artificial Intelligence Study," 2022

Agentic AI and GenAI

Both technologies are interrelated, but have significantly different capabilities. GenAI studies patterns in very large data sets and outputs human-like responses to the prompts it receives. GenAI isn't autonomous, lacks decision-making capabilities (beyond pattern recognition), and isn't goal oriented. Many AI agents incorporate GenAI in workflows to communicate and can be heavily reenforced with enterprise-specific information to add value to its outputs. AI agents that communicate with customers (humans) as part of its goal-oriented task often leverage GenAI as the enabling technology.

Agentic AI Adoption Trends

While the topic of AI in general has been growing in importance since the launch of OpenAI's ChatGPT over two years ago, successful adoption of AI across the enterprise has been curtailed by issues like data quality and availability, as well as the need for highly trained data science skill sets. This has limited adoption for many companies and stifled their ability to scale successfully.

To adopt and execute AI projects, companies need to demonstrate maturity in the following seven areas:

- Identify AI leadership in the organization
- Security standards and compliance processes
- Data access/use policies
- Legal considerations and impact
- AI program standards, policies, and ethics
- Project governance oversight and frameworks
- Enterprise architecture requirements for AI

BARC research shows that only 21% of 370 global respondents[2] have fully operationalized each of these requirements as part of an AI strategy foundation.

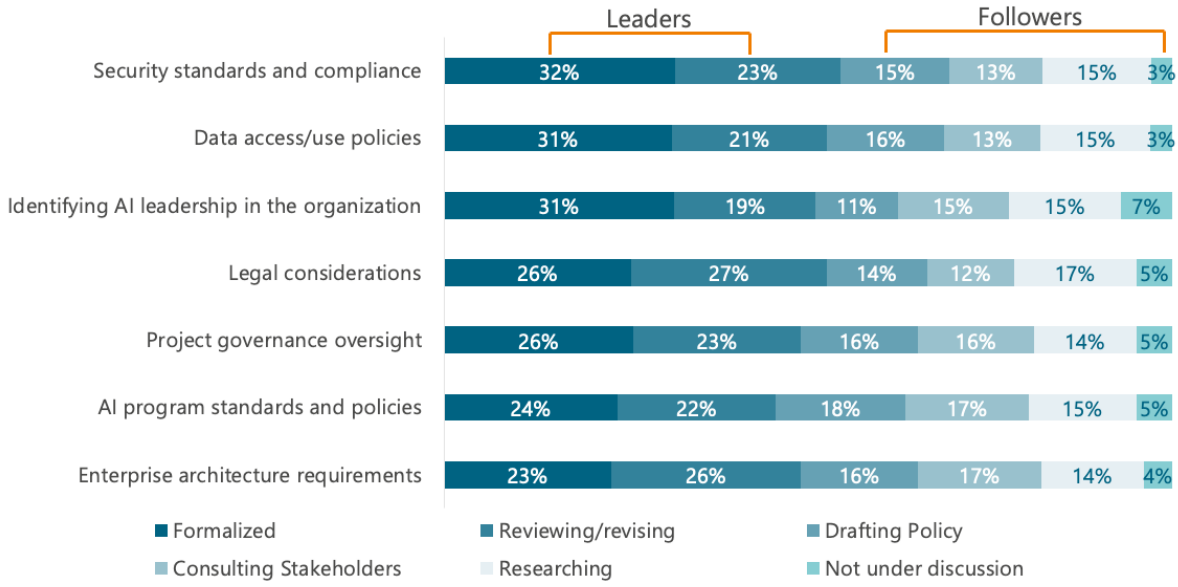

[2] BARC Study: Preparing and Delivering Data for AI Adoption Trends, Requirements and Best Practices, March 2025

Addressing these complex organizational and policy issues generally requires firms to rely on outside experts to help guide and deliver these programs. This foundation is equally important for companies leveraging agentic AI.

BARC Research on AI Agent Adoption

Specifically looking at AI agent adoption, it's easy to see that companies are all in on this AI technology. Of BARC survey respondents, 32% have agents in limited or full production, and 26% are in pilot or experimentation stages[3].

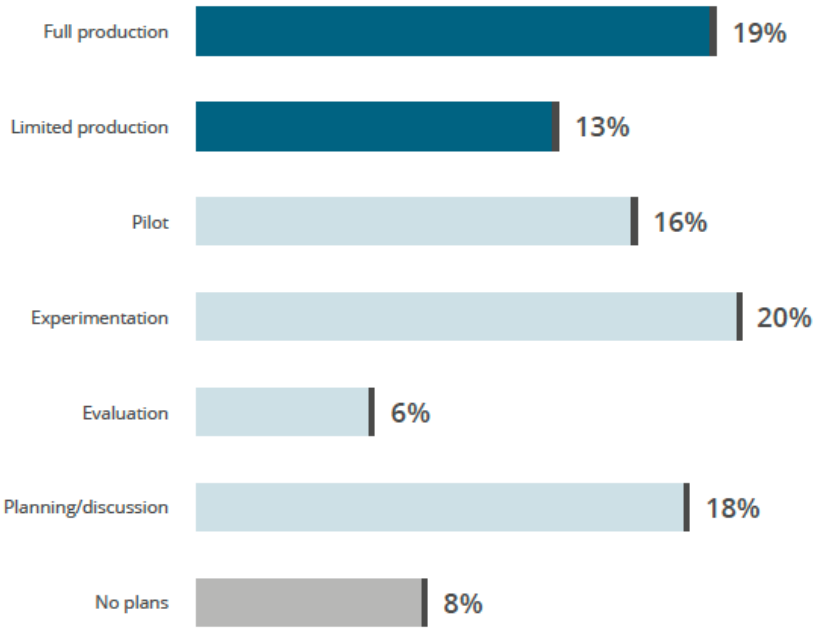

Figure 1: Overall, please describe how your organization is adopting AI agents / autonomous AI systems. (n=285)

Delivering Value with Agentic AI

Fast Track to AI Value

Agentic AI is a quicker, faster route to AI value, scale, and cost savings than many traditional AI competencies. Its autonomous capabilities, coupled with deep analytic insights and goal-defined actions, deliver a full circle value chain not generally found in other types of applications. Agents allow for tactical execution at scale and speed workflows to operate continuously, never requiring a break or time off. The volume of work alone creates a compelling cost-saving scenario for companies. Agents, once deployed, require limited manual input, oversight, and free up humans to contribute in new and creative ways. Additionally, agents can be reused and redeployed reducing initial development costs.

[3] BARC Study: Preparing and Delivering Data for AI Adoption Trends, Requirements and Best Practices, March 2025

Insight to Action Automation

Sophisticated analytic systems have benefited from machine learning-based technologies for many years, empowering companies to move beyond the traditional four types of analytics:

1. Descriptive: What happened in the business?
2. Diagnostic: Why did it happen?
3. Predictive: What's likely to happen?
4. Prescriptive: What needs to happen or change?

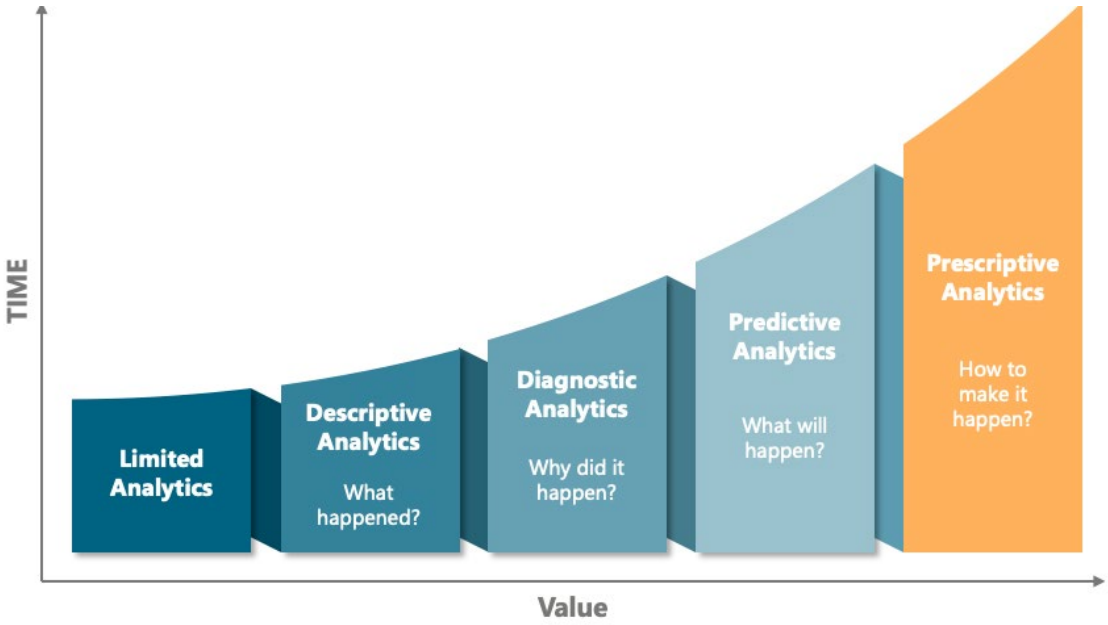

Agentic AI delivers the next and perhaps most valuable step – action. Agent-driven applications can leverage all these analytics types and complete the value cycle by acting on the parameters of the goals provided to the agent.

Single and Multi-Agent Systems

As agentic AI arrived on the scene in 2024, single-agent, single-task solutions were quickly eclipsed by high-value, complex, multi-agent solutions. Teams of agents, each with its highly focused capabilities can be combined to learn from each other and train each other to provide the best overall output and action towards the parent agents' overall defined goal. Simple agents will continue to exist in the marketplace. The adoption chart above shows that 32% of study respondents have deployed an agent, but it's likely these are simple, single-agent workflows. The deepest value will come from more sophisticated, multi-agent applications.

Agentic marketplaces are growing quickly, providing companies with easy access shopping for prebuilt and industry-specific agents that can be imported to add value and performance to multi-agent solutions. These marketplaces are empowering smaller companies to get involved with agentic AI quicker than building everything independently. Marketplaces are a great resource for companies looking to leverage domain and industry-specific agent functionality. Sites like AI Agent Directory list more than 1,200 AI agents, ranging from a medical scribe agent designed to convert text to medical notes to a quant agent to analyze the crypto market.

As an agent leverages more functionality, they can share data and insights and help create a more reliable workflow. This is often done via agent reflection, a prompting strategy used on LLMs in the workflow to have the LLM reflect on and critique its own outputs. Often, an agent in the workflow will add additional data or information that enhances context with the model providing better input to the next step in the workflow.

Additional agent examples can include compliance, ethics, and corporate policy agents that keep processes in check and report back to enterprise systems for better transparency and governance. Autonomous decision-making frameworks referenced in the following section play a necessary role in enabling all agent workflows and is a required element for success.

Technology Considerations

The Core Technologies Driving Agentic AI

Language Models (LLMs, SLMs and DLMs)

In the early days (more than two years ago), LLMs fueled GenAI outputs. These models are huge and contain an extremely wide set of information. The cost of building and running these models is prohibitive and very few companies were successful in building out their own – that work has been left to the hyperscalers who leverage billions of pieces of data to train their models.

A new trend serves agentic AI through small language models (SLMs) and domain language models (DLMs). These highly specific, lighter-weight models run faster and at lower cost. Agent builders are leveraging SLMs and DLMs to drive better context and learning in agentic AI applications.

Companies can adopt a multi-model approach to enable better accuracy in predictive workflows. By creating an ensemble model from the output of multiple machine learning models, end users can produce an extremely accurate prediction model, making an agent more accurate. Smaller sized models will also impact overall AI costs and perform faster, making them a valuable part of an enterprise AI strategy.

Domain language models (DLMs): These highly specific, lighter weight models run faster and at lower cost.

Advanced Computing Hardware

Graphics processing units (GPUs) are a critical part of AI strategies. These chips are especially proficient at taking on the compute demands for machine learning, scientific simulations, and AI/agentic workloads. Nvidia, AMD, and Intel are leaders in this market. Hyperscalers AWS, Google, and Microsoft, along with Meta and Oracle, are all hosting 50,000 to 150,000 GPUs within their cloud offerings to meet the growing needs of AI and agentic AI workloads. These investments position hyperscalers at the center of most AI strategies, as it is financially prohibitive for most companies to build this infrastructure independently.

Model Context Protocol (MCP)

MCP is quickly becoming a universal tool standard for agentic AI. MCP reduces the time needed for complex tool integrations by managing access to external services, allowing agents to reliably invoke and interact with those tools in dynamic workflows. MCP services enable agents to interact with structured prompts, external functions, tools, and real-world data sources – making them more capable and autonomous. Many tools are now accessible through platforms that support agent-based architectures. MCP can eliminate custom coding for each agent and helps to overcome incompatible AI frameworks. As agentic AI becomes more complex, companies will need this type of interface to save time and money, and avoid agent chaos.

Autonomous Decision-Making Framework

Autonomous decision-making frameworks are a required system for agent creation. These systems help design agents "think" and act strategically to carry out specified goals, while reacting to planned and unplanned scenarios and information:

- **Quantifiable Goal Design:** Goals need to be very specific, measurable, achievable, relevant, and time-bound (SMART). These goals should include milestones, and an execution hierarchy to lead to successful completion and outcomes. Clarity of purpose is paramount in agent design.
- **Learning and Improvement:** For agents to perform at a high level, knowledge accumulation (learning) needs to be programed into the agent workflows and processes. By providing continuous feedback and updating the models via machine learning integration, agents can perform more accurately.
- **Human in the Process:** Integration with humans based on speed of the application, risk, complexity, and governance is essential.
- **Integration to Tools and Systems:** Simple access to data, information, policies, governance and privacy systems, enterprise marketing, customer and service systems, and API and data connectors, all need to be available in this framework.
- **Evaluation Process:** A process to determine performance based on the measurable and relevant goal needs to be in place, as well as alerts to loop humans into failing processes, and regular evaluations and testing.

These frameworks enable autonomous agents to stay on the right path, learn, improve, and execute. Without them agentic AI can be unmanageable and risky.

Integration to Business Intelligence and Data Platforms

Agentic technology is being built into leading business intelligence (BI) and data platforms, along with GenAI capabilities. Many of these platforms are leveraging their connectivity to enterprise data and metadata to deliver agent building and creation environments. Common agents in these environments focus on the data, creating AI-driven data definitions, mapping, and connecting data via integration paths and APIs – fairly simple work agent technology, but a valuable time saver.

Data is the Oxygen for AI

Agentic AI cannot be built on a soft data foundation. Strong data management, access, and orchestration drive value and accuracy for agentic AI. Most companies gravitate to familiar types of data initially and BARC research shows that 69% of respondents are in production or POC using structured data. However, 53% of respondents have done the same with image, video, and sound data and 54% have accomplished it with unstructured data and documents.

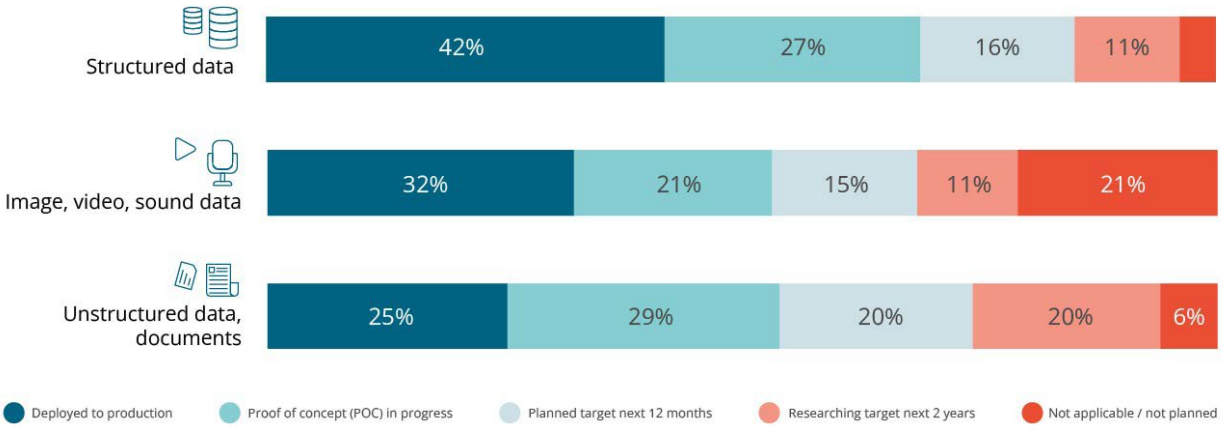

Figure 2: How is your organization utilizing each of the following data types for AI? (n=324)

Unstructured data and documents play a critical role in retrieval-augmented generation (RAG) workflows that help address custom use cases, while mitigating the risk of model hallucinations within agentic workflows.

Responsible Standards for Agentic AI

Responsible AI

Compliance and regulatory standards impact companies from an "outside-in" perspective. Smart organizations align policies to engage with these standards. Responsible AI is most often an internal set of guidelines and polices designed to align to corporate strategy and values. Responsible AI acts as a corporate guard rail to avoid over-innovation, protect data, and define policies for accuracy and action with agentic AI.

BARC Research Insights

Recent research continues to highlight the areas of responsible AI that are being prioritized by companies. Data privacy, security of AI systems and models, along with transparency and auditability are excellent places to focus. It's concerning that model bias, job impact, and data bias rank so low and is especially problematic for companies planning to leverage agents across the enterprise. This is a mistake that all forms of AI will amplify quickly and can cause significant issues with agent performance and actions.

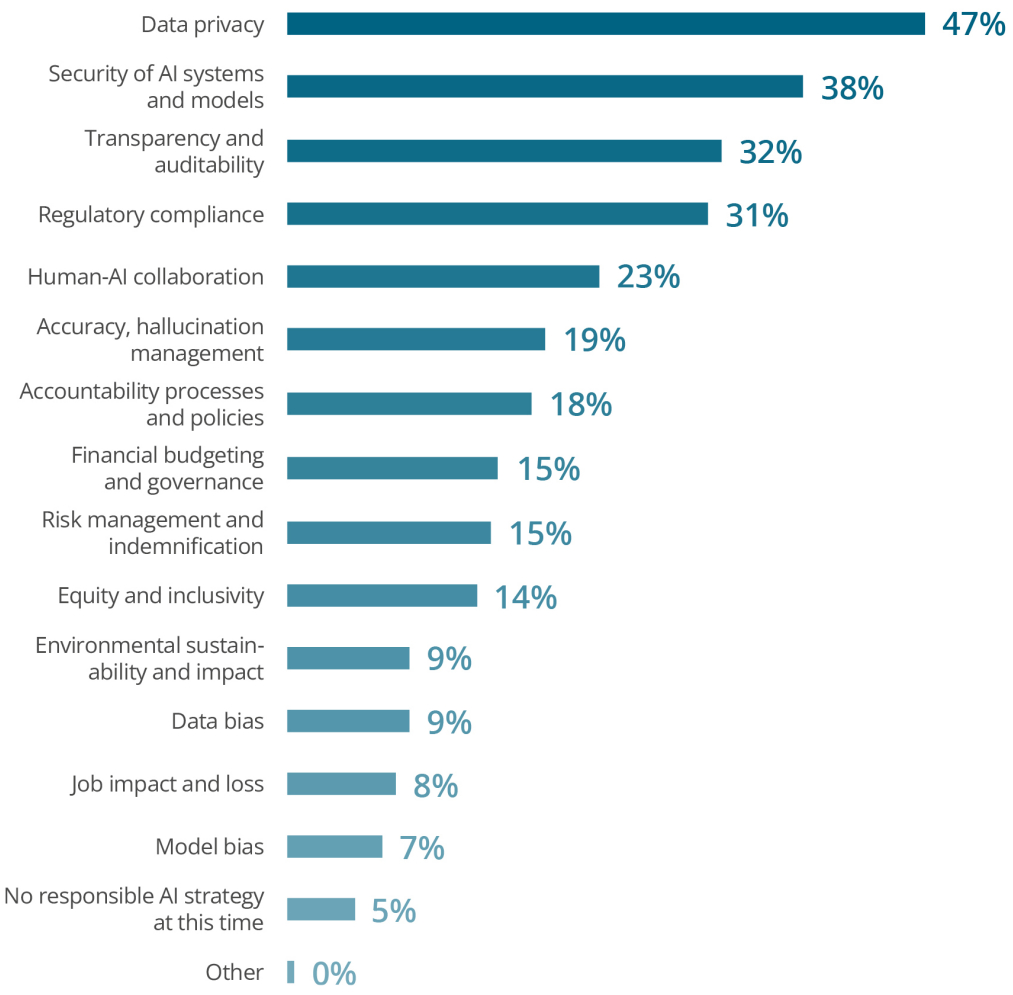

Figure 3: What components of Responsible AI is your organization prioritizing? (n=335)

Overcoming Agentic AI Challenges

A guiding principle for AI and agentic AI specifically is to avoid over-innovation. Often, the buzz and excitement about new technology opens the door for shadow IT projects that operate outside of the guardrails of corporate IT, causing significant tension and sometimes, costly errors. The greater the autonomy of an agent, the higher the risk and, often, the higher the reward. This balance is critical to understand and it's important that any AI-based initiative aligns with corporate governance, security, and privacy guidelines.

Robust Governance Frameworks

Governance is different for every company; as agentic AI spreads throughout your enterprise, it immediately becomes a non-negotiable part of the AI strategy. Governance frameworks become the foundation for policy, ethics, and compliance agents that are included in multi-agent processes (agents governing agents) to keep complex workflows on target. The following are critical parts of a strong governance framework:

- **Transparency:** Clear information on what agents should and should not do, data sources, decision logic, and processes – much like a coding project
- **Risk Management:** Understanding potential failure and ramifications, alignment with corporate rules and regulations
- **Ethical Standards:** Alignment to corporate ethics standards, data and model bias, and fairness guidelines
- **Accountability:** Oversight and resolution procedures, review boards specifically assigned to corporate AI practices
- **Ongoing Review:** Agile processes to track and review agent performance and its alignment to corporate guidelines, as well as legal and regulatory topics

Robust Security and Model Governance

Security and model governance needs to be prioritized when implementing agentic AI. BARC Research shows that security of AI systems and models is the top concern for companies. Of all respondents, 40% ranked it first, followed by human-AI collaboration and data privacy. Agents need to be slipstreamed into existing security and governance platforms to avoid making significant mistakes with agent technology.

Humans bring emotional intelligence, nuanced understanding, and creativity skills to problems.

Humans in the Process

Autonomous action with AI agents delivers value, but as stated above, the greater the autonomy, the greater the risk. Part of a successful agentic strategy is understanding how and when to keep humans in the loop. Humans bring emotional intelligence, nuanced understanding, and creativity to problems. Agents lack these skill sets making it critical to include humans at the right time with each process.

Balance your strategy considering the speed at which the agent is completing its process (goal), how complex the process is (single agent, multi-agent), the overall risk of agent performance, any regulatory and compliance issues, and the impact of mistakes. Always wrap agents in enterprise governance and security layers.

	Humans out of process **Autonomous**	**Humans engaged** **As Required**	**Humans required** **Engaged**
Agent Speed	Sub second	>3 minutes	As designed
Complexity	Low → High		
Risk	Low → High		
Governance	Yes	Yes	Yes

Industry Use Cases

In the following examples of agentic AI in action, an insurer and a pharmaceutical company each explored how multi-agent systems could support their teams in areas such as quality assurance and clinical code validation – domains where accuracy, scale, and regulatory oversight are critically important.

Endava supported both initiatives, working closely with each company to shape the right solutions for their business needs. Through its partnerships with world-leading organizations, Endava supports the adoption of next-generation technology, designing and delivering AI-native solutions that accelerate operational efficiency, ensure compliance, and create lasting value.

Insurance Claims

When an organization reviews thousands of medical assessment documents each week, a manual quality-assurance queue drains time, hides errors, and exposes you to compliance liabilities. That was the case for a leading claims-validation firm serving insurers and attorneys. Specialist teams had to spend hours checking each file, tracking progress in spreadsheets, and re-routing mistakes. Delays rippled through claims decisions and left managers guessing about case status.

Endava replaced that bottleneck with a set of AI agents, each tuned to a single QA task. A Field-Validation Agent compared every predefined data field to source records. A Language Agent applied U.S. English rules to spot spelling and grammar issues. A Medical-Response Agent read clinical notes and confirmed that doctors answered all required questions. These agents run inside a secure web app powered by Claude Opus and Azure Cognitive Services, returning results quickly and logging every step in a real-time dashboard.

The impact of the AI agents was immediate:

- **Processing time**: cut from hours to under 5-10 minutes per case
- **Staff efficiency**: freed analysts for higher-value reviews, avoiding new hires
- **Visibility**: live status tracking replaced ad-hoc emails and spreadsheets
- **Compliance**: automated checks reduced HIPAA exposure, costly reworks, and secured the handling of sensitive data with regulatory safeguards

With scalable cloud services, the firm can deploy more agents as volume grows or regulations change. By delegating repeatable verification to specialized AI agents, you keep human experts in the loop and focused on judgment calls, shorten claims cycles, and raise confidence in every medical assessment you deliver.

Pharmaceutical, Automating Clinical Code

A top 10 pharmaceutical company, following standard industry practices, employs a double programming approach for creating clinical code to reduce the risk of error, though this increases costs and slows time to market. Regulatory approval for clinical trials demands strict adherence to Clinical Data Interchange Standards Consortium (CDISC) and analysis data model (ADAM) standards. Manually ensuring compliance was slow, error-prone, and risked regulatory delays. Lengthy coding and verification cycles can delay approvals, leading to millions in costs. These processes often depend on expensive and

scarce specialized clinical developers, further hindering timely patient access to critical therapies. Additionally, regulatory requirements add complexity, extending review timelines.

To address these challenges, the company sought an AI-driven solution that would accelerate code creation and reduce clinical trial costs and timelines while maintaining high levels of accuracy and ensuring compliance.

To address these challenges, Endava deployed Morpheus, Endava's multi-agent AI platform, leveraging a reflection pattern where AI-generated code was continuously refined by reviewer agents.

These AI agents ensured:

- Automated Code Generation: generating R, Python, and SAS clinical trial code
- Regulatory Compliance Validation: AI reviewer agents enforced CDISC and ADAM standards, running automated unit tests to minimize errors and ensure submission readiness
- Scalability Across Studies: The solution was designed to be adaptable, enabling deployment across multiple clinical studies and therapeutic areas

The solution utilized:

- Generator agent: writes draft code from study specs and a knowledge graph
- Reviewer agents: runs unit tests, compares outputs with CDISC and ADAM templates, and refines code until every rule passes
- Audit layer: records each action in a dashboard for inspectors
- Scalability: one click rolls the agents to new studies and therapeutic areas

Business impact

- Projected $35 million annual reduction in contract-programmer spend
- Dataset build time cut from days to minutes
- Performance proven across eight identified studies
- Full traceability gave auditors immediate confidence in every submission

The Future of Agentic AI

The Role of Agents in the Enterprise

Agentic AI will grow exponentially across the enterprise in coming years. As referenced above, hundreds to thousands of agents will operate in every corner of the enterprise. Over the next five years, agents will become the default resource for employees trying to solve challenges and impact company KPIs.

Emerging Technology and Trends

- Agents using tools will become commonplace: NLP, database connectors, APIs, language translators, and image recognition have already impacted complex agent capabilities. Third-party resources, data, and capabilities will be folded into agents to add context, understanding, and drive actions.
- Agents as teammates: As employees, we'll work alongside agents in the workplace, incorporating its work with ours to make faster decisions, and shed tactical work to allow humans to make a great impact with their time.
- Agentic AI will leverage a multi-model world: Small and domain-specific models will play a key role in helping agentic AI execute faster and more accurately. Of BARC survey respondents, 40% include financial models; 39% prioritize coding models, and 25% are leveraging science-based specialty models in their AI environment.[4] These highly focused domain models increase accuracy and help agents operate with more context and accuracy.

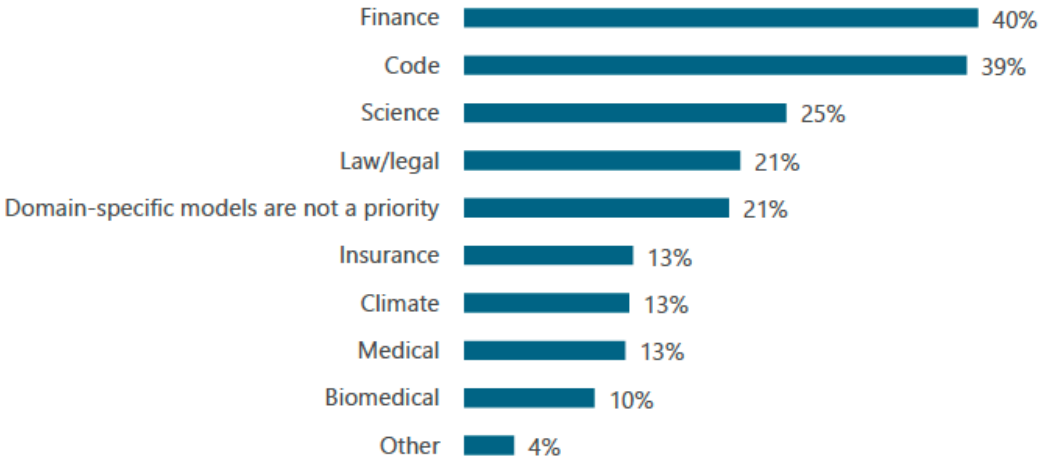

Figure 4: Which Industry or domain-specific models are included in your AI model strategy? (n=335)

Autonomous Teams

Companies are already hiring AI agents in a similar fashion as human employees. If your organization is exploring how to augment your teams with agents, there are many things to consider. Companies who are weighing the cost of operating a team of agents versus traditional headcount or benefit costs will need to follow a few commonsense guidelines:

- **Performance Metrics and Agent Evaluation:** Performance criteria and the ability to monitor performance in real time will be a requirement. A governing agent will likely be used in these scenarios.
- **Public Relations and Reputation Management:** Plan for and be ready to defend the decision to hire agents instead of people. Proactive communications internally will be extremely important along with external messaging on the topic.

[4] BARC Research Optimizing Your Architecture for AI Innovation, March 2024 Merv Adrian, Shawn Rogers

- **Workforce Integration:** Human and AI collaboration is a new concept. Collect feedback and focus on agents adding value to its human counterparts and highlight these work culture benefits. Clearly define how humans will interact and leverage agents in the workplace.

There are many more considerations for companies adding agents to their workforce – remember to communicate authentically with your teams and be transparent. Building autonomous teams can have significant cost benefits and can offload tactical work from humans to free them up to leverage their skills in a more meaningful way within the business.

Recommendations

The impact of agentic AI on enterprise businesses is undeniable. Leveraging the technology is an opportunity that can't be ignored. This technology is moving fast, but as shared earlier, only 21% of research respondents have operationalized all aspects of their foundation for AI and agentic AI.

Where to begin:

1. **Build the foundation for success.** Operationalize these seven critical competencies before jumping into agentic AI: security standards and compliance processes; data access/use policies; identify AI leadership in the organization; legal considerations and impact; project governance oversight and frameworks; AI program standards, policies, and ethics; and enterprise architec-ture requirements for AI.
2. **Data is the oxygen for AI.** Don't ignore data quality, access, and orchestration challenges. 47% of BARC research respondents list data quality as the top concern when measuring AI success. AI and agentic AI are tremendous forcing factors to address data quality issues to succeed.
3. **Don't do agentic AI alone.** Partner with experienced firms who can de-risk your AI innovation strategy. Experienced partners provide guidance that delivers a path to avoiding common and costly mistakes, while enabling quicker timelines and greater return on investment.

About BARC

BARC is the leading analyst firm for data & analytics, AI, corporate performance management (CPM) and ESG with a reputation for unbiased and trusted advice. Our expert analysts deliver a wide range of research, events and consulting services for the data & analytics community. Our innovative research evaluates software, vendors and service providers rigorously and highlights market trends, delivering insights that enable our customers to innovate with data, analytics and AI. BARC's 25 years of experience with data strategy & culture, data architecture, organization and software selection helps clients transform into truly data-driven organizations.

Research

BARC user surveys, software evaluations and analyst advisory services along with expert driven content such as research notes, trend analysis and blogs give organizations the confidence to make the right decisions. Our independent research gets to the heart of market developments, evaluates software, vendors and service providers thoroughly and gives valuable ideas on how to turn data, analytics and AI into added value and successfully transform businesses.

Consulting

The BARC consulting practice is entirely focused on translating companies' requirements into future-proof decisions. The holistic advice we provide helps companies successfully implement their data & analytics strategy and culture as well as their architecture and technology.

BARC's research and experience-founded expert input sets organizations on the road to the successful use of data & analytics, from strategy to optimized data-driven business processes.

Events

At BARC events, leading minds and industry experts come together to share insights and drive innovation. Our conferences, roundtables and online webinars attract over 10,000 participants annually, offering a unique blend of information, inspiration and interactivity. These events provide a platform to exchange ideas with peers, explore emerging trends and gain expert perspectives on market developments.

By engaging with thought leaders and industry practitioners, participants discover actionable strategies to enhance their business and stay ahead in the evolving world of data & analytics.

BARC

About Endava

Global Headquarter
125 Old Broad Street,
London, EC2N 1AR
United Kingdom

+44 20 7367 1000
www.endava.com

Endava is a leading provider of next-generation technology services, dedicated to enabling its customers to accelerate growth, tackle complex challenges and thrive in evolving markets. By combining innovative technologies and deep industry expertise with an AI-native approach, Endava consults and partners with customers to create solutions that drive transformation, augment intelligence and deliver lasting impact. From ideation to production, it supports customers with tailor-made solutions at every stage of their digital transformation, regardless of industry, region or scale.

Endava's clients span payments, insurance, finance and banking, technology, media, telecommunications, healthcare and life sciences, mobility, retail and consumer goods and more. As of December 31, 2024, 11,668 Endavans have helped clients break new ground across locations in Europe, the Americas, Asia Pacific and the Middle East.

Our experience with intelligent agents and automation enables us to support clients as they introduce systems that can act independently within defined parameters. We focus on making these technologies practical and secure, helping organisations apply agentic AI to real business challenges in a way that delivers measurable results.

The views, opinions, and recommendations expressed in this report are those of the author (BARC) and do not necessarily reflect the official policy, position, or endorsement of Endava. Any reliance placed on the information in this report is strictly at the reader's discretion.

Data Decisions. Built on BARC.

www.barc.com

United States
BARC US
13463 Falls Drive
Broomfield, CO 80020
USA

Germany
BARC GmbH
Berliner Platz 7
D-97080 Würzburg
+49 931 880651-0

Austria
BARC GmbH
Hirschstettner Straße 19 / I / IS314
A-1220 Wien
+43 660 6366870

Switzerland
BARC Switzerland GmbH
Buchhaldenstraße 7
CH-5442 Fislisbach
+41 76 340 35 16