CORPORATE PERFORMANCE IMPROVEMENT

Making Agentic AI Work:

Technology Foundations, Operating Model Implications, and Evaluation Criteria

The primary determinant of success in agentic AI is not experimentation speed, but the technology foundation chosen to support scale and control.

Executive Summary

Agentic AI is emerging as a transformative capability, enabling organizations to move beyond task-based generative AI toward autonomous, goal-driven systems that orchestrate end-to-end process automation. This evolution has the potential to unlock significant operational efficiencies and strategic differentiation, but success depends heavily on selecting the right technology foundation.

The agentic AI market is already entering a proliferation phase. The current landscape is highly fragmented, with startups and major cloud providers offering overlapping solutions. Prior technology waves show that proliferation often leads to overspending, integration challenges, and vendor instability. Without a disciplined approach to technology selection, organizations risk fragmented architectures that are difficult to scale or govern.

To avoid these pitfalls, organizations must take a structured, pragmatic approach to agentic AI technology selection grounded in enterprise requirements rather than vendor hype.

Key insights include:

- **Enterprise-critical capabilities:** Several new technologies are essential for deploying agentic AI at scale, including tools for runtime security and governance, observability, tool and API management, contextual memory, data protection, and human-in-the-loop oversight.
- **Technology foundation options:** Four primary approaches exist—frameworks, platforms, domain-specific solutions, and hybrid models—each with distinct strengths and trade-offs.
- **Operating model alignment:** Technology choices directly influence the organizational structures required for supporting them, including centralized Centers of Excellence for frameworks, democratized development models for platforms, and hybrid approaches for large enterprises.
- **Selection criteria:** Effective agentic AI technology selection decisions balance strategic priorities, technical requirements, financial considerations, and operational feasibility.

Organizations should prioritize high-value use cases, assess readiness honestly, and select technology foundations that align with long-term goals. Enterprises that act decisively now will be better positioned to scale, govern, and realize value from agentic AI.

> Agentic AI is moving rapidly from experimentation to operational expectation.

Introduction

Agentic AI is fast moving from experimentation to operational relevance. While popular generative AI tools improve productivity at the individual task level, agentic AI drives enterprise-scale transformation. Organizations are widely investing in autonomous agents to unlock new efficiencies, insights, and competitive advantages across a wide range of complex workflows, but realizing value requires more than enthusiasm or isolated pilots. It demands a deliberate approach to technology selection aligned to enterprise realities.

Many challenges associated with scaling agentic AI solutions—security concerns, unclear accountability, integration complexity, and stalled deployments—are not inherent to the technology itself. They are consequences of choosing technology stacks that only meet the functional requirements of use cases but do not align with team skills, governance, or operating models.

> The agentic AI market is already entering a proliferation phase characterized by overlapping tools, unclear ownership, and escalating complexity.

In this environment, the challenge is no longer whether agentic AI is viable, but how organizations cut through market noise to make technology choices that will hold up under enterprise-scale demands. This makes disciplined technology selection—not experimentation alone—the critical differentiator.

> Organizations that move too quickly without a selection framework risk being locked into tools that do not scale with enterprise needs.

Navigating the Proliferation of Agentic AI Technology

Agentic AI is entering a familiar phase of rapid market expansion and intense competition. Foundation models have made "intelligence" inexpensive and widely accessible, enabling a surge of new vendors and solutions for nearly every niche and use case. The result is a crowded landscape:

- Startups offering highly specialized solutions with narrowly focused agent capabilities
- Large enterprise software vendors delivering broad, infrastructure-heavy platforms that are generic, complex, and full of gaps

History suggests this phrase brings predictable, significant challenges. Organizations end up overspending on tools without realizing their full potential, duplicating capabilities across overlapping platforms, and enduring lengthy implementation and integration cycles that consume technical team capacity. Then, whether through consolidation or failure, vendor instability often leaves organizations with unsupported tools and costly replacement projects.

Given these trends, a consolidation phase is inevitable. Larger players will either acquire niche innovators or replicate their capabilities to render them obsolete, leaving early adopters to unwind fragmented architectures.

Organizations that treat agentic AI technology selection as a procurement exercise rather than an operating-model decision risk repeating the same failures seen in prior technology waves.

From Market Noise to Enterprise Reality

Agentic AI technologies do not operate in isolation. They embed deeply into workflows, systems, and decision-making processes. As a result, the chosen technology foundation must integrate seamlessly into the broader context of the organization and align with:

- Governance and compliance requirements
- Data and security policies
- Talent availability and engineering maturity
- Desired speed of deployment vs. depth of customization

> Agentic AI tool selection is not a procurement exercise. It is an operating-model decision.

For example:

- Framework-based approaches prioritize deep customization and strategic differentiation, but require strong governance and skilled teams
- Platform-based approaches support rapid deployment and democratized development, but must fit into existing compliance and data management structures
- Hybrid approaches balance agility and control, enabling scale across diverse use cases without sacrificing oversight

Translating organizational realities into effective technology decisions requires clarity on what agentic AI must be able to do at scale. Before evaluating specific tools or vendors, organizations must first understand the enterprise capabilities required to operate autonomous agents safely and reliably.

> Technology decisions made early in the agentic AI journey have long-term implications for governance, cost, and scalability.

What Capabilities Are Included In An Agentic AI Foundation

Early generative AI tools have delivered clear productivity gains by assisting individuals with discrete tasks such as content creation, summarization, and analysis. These tools, however, remain largely prompt-driven and human-supervised, limiting their ability to automate end-to-end processes or own outcomes at an enterprise level.

Agentic AI represents a shift from task execution to outcome ownership. Agents can plan actions, coordinate across systems, execute multi-step workflows, and adapt based on results. This autonomy introduces new operational, security, and governance requirements that technology stacks were not specifically designed to address.

To deploy agentic AI safely and effectively at scale, organizations need a technology foundation that embeds governance, security, observability, and control directly into how agents operate, not as afterthoughts.

Core capabilities include:

Runtime security and governance
Agentic AI systems must enforce organizational compliance, security, and governance policies at runtime. Built-in safeguards, guardrails, access controls, and privilege management help ensure agents operate within approved boundaries and adhere to regulatory requirements.

Observability and tracing
Comprehensive monitoring and end-to-end visibility into agent behavior are critical for reliability and trust. Organizations need insight into agent performance, reasoning steps, and tool usage to detect anomalies, validate outcomes, troubleshoot issues, and maintain quality across automated workflows.

Tool and API management
Centralized management of tools and integrations with permission controls enables reuse, prevents duplication, and limits uncontrolled proliferation. Permissioned access ensures agents use approved resources while maintaining security and operational consistency.

Contextual memory and retrieval
Agents require both short- and long-term memory to reason effectively across tasks. Short-term memory supports individual conversations and process execution, while long-term memory supports ongoing learning of organizational context. Governed memory improves decision-making while reducing the risk of hallucinated or outdated context.

Data protection and privacy middleware
Protecting sensitive information is non-negotiable. Middleware provides real-time handling of personally identifiable information (PII), masking, and policy-driven transformations to ensure compliance with privacy regulations.

Human-in-the-loop (HITL) supervision
Even in autonomous systems, certain critical actions require human review and approval. With traceability and approval workflows, HITL interfaces enable oversight and control for high-risk decisions without sacrificing automation benefits.

The goal is to establish a technology foundation that incorporates all these capabilities. Not all must be implemented on day one, but the chosen technology foundation must support their introduction as use cases mature.

Four Technology Foundations for Agentic AI

Selecting the right technology foundation is one of the most critical decisions for organizations embarking on their agentic AI journey. The choice depends on factors such as technical expertise, customization needs, compliance requirements, and alignment with existing ecosystems.

Broadly, there are four options for building the technolog foundation for agentic AI: **frameworks, platforms, domain-specific solutions**, and **hybrid models**. Frameworks offer maximum control and customization, while platforms offer speed and scalability. Domain-specific solutions provide industry-focused customization with compliance, while hybrid models seek to balance both agility and operational stability.

Each foundation offers distinct advantages and trade-offs and is suited to different organizational profiles, noted in the table below.

Technology foundation	Description	Best for	Strengths	Trade-offs	Examples
Frameworks	Libraries—often open-source—for building custom agent architectures	Organizations with strong engineering capabilities and a need for customization or desire for vendor independence	Maximum control, intellectual property (IP) ownership, flexibility	Higher build and maintenance cost	LangGraph, CrewAI, Microsoft Agent Framework
Platforms	Managed services that provide agent infrastructure, orchestration, governance, and scale	Organizations prioritizing speed to market, operational efficiency, and simple out-of-the-box solutions integrated with existing data or cloud ecosystems	Rapid deployment, minimal engineering talent overhead, scalability, built-in compliance and controls	Architectural constraints, reduced flexibility, potential vendor lock-in	Google Vertex AI, AWS Bedrock, Databricks, Microsoft Foundry
Domain-specific solutions	Pre-configured, pre-trained, pre-governed agents tailored to specific industries or functions	Highly regulated or mission-critical use cases	Embedded compliance, deep system integration	Limited flexibility or applicability beyond target domain	Salesforce Agentforce, Hippocratic AI, Harvey AI
Hybrid models	A combination of platforms and frameworks aligned to different use cases	Large enterprises with diverse needs, mixed technical maturity, and desire for both stability and agility	Balance of control and speed	Increased architectural complexity	--

How Technology Foundations Shape the AI Operating Model

Selecting a technology foundation not only involves a technical decision, but also directly shapes how agentic AI is owned, governed, and delivered across the organization. As a result, each technology approach naturally aligns to a different AI operating model, each with distinct benefits, risks, and resources requirements.

The table below outlines the three primary AI operating models—**Center of Excellence (COE), democratized development**, and **hybrid**—and how they align with agentic AI technology foundations.

Operating model	Description	Key benefits	Technology alignment	When this model works best
Center of Excellence (COE)	A centralized team of AI experts responsible for designing, building, governing, and maintaining agentic AI solutions across the enterprise	• Strong governance, security, architectural consistency, and **maximum control** • **Reusable components**, shared libraries, and standardized deployment pipelines • Reduced risk through centralized expertise and oversight	• Best aligned with **framework-based approaches** (e.g., LangGraph, Microsoft Agent Framework) • Enables development of custom orchestration layers, reusable agent libraries, and centralized observability platforms	• Organizations with **mature engineering and data teams** • Strong regulatory or compliance requirements • Strategic differentiation is a priority (e.g., financial services, healthcare, regulated energy)
Democratized development	A distributed model in which business users and citizen developers build agents using low-code or no-code tools under defined governance guardrails	• Faster speed-to-value by enabling non-technical teams to create solutions • Reduced dependence on central IT for routine use cases • Increased experimentation and innovation at the business-unit level	• Requires **platform-based solutions** (e.g., Google Vertex AI, Microsoft Foundry) • Visual interfaces, pre-built templates, and automated compliance controls are critical • Strong built-in guardrails, audit trails, and approval workflows are essential to minimize risks	• Organizations facing **talent gaps** or urgent delivery timelines • Use cases focused on operational efficiency rather than differentiation • Environments where governance and training programs are well established to avoid low-quality outputs and compliance violations
Hybrid model	Combines centralized governance with distributed execution. A COE sets governance, standards, templates, and platforms, while business units configure, customize, and deploy agents within approved boundaries	• Balance of **control and agility** • Scales delivery without creating central bottlenecks • Enables innovation while maintaining enterprise standards	• Typically combines **managed platforms with governance layers** • Requires multi-tenancy, API extensibility, granular access controls, and centralized monitoring • Often integrates resuable templates and frameworks for specialized or high-risk use cases	• Large enterprises with diverse use cases across functions or regions • Mixed technical maturity across teams • Need to scale agentic AI without sacrificing compliance or oversight

Technology selection and operating model design must work hand-in-hand to achieve maximum impact quickly. Misalignment between the two is a common cause of stalled or fragmented agentic AI initiatives. Successful alignment promotes scalability, compliance, and speed to value – all critical for realizing the full value of agentic AI.

Criteria for Technology Foundation Selection

With technology options and operating models clearly defined, the final step is determining which combination best fits an organization's specific context. Selecting the right agentic AI foundation requires a disciplined balancing of competing priorities, as no single approach is universally "best". This requires an evaluation framework that weighs strategic intent against technical feasibility, financial sustainability, and operational readiness.

A robust evaluation considers four core dimensions—strategic, technical, financial, and operational. Each dimension surfaces different risks and constraints, and their relative importance will vary by organization. The objective is not to optimize any one dimension in isolation, but to arrive at a balanced decision aligned with enterprise realities.

Strategic

Strategic evaluation focuses on how technology choices align with competitive positioning and long-term business priorities. In most organizations, this dimension carries the greatest weight.

Key considerations include:

- **Differentiation vs. commodity capabilities**
  Capabilities that provide meaningful strategic differentiation—such as proprietary algorithms, unique workflows, or novel ways of embedding intelligence into core processes—often justify custom builds using frameworks. In contrast, commodity use cases (e.g., standard customer service chatbots or internal productivity assistants) are typically better served by platforms that prioritize speed and cost efficiency.
- **Control and IP requirements**
  Organizations that view AI as a core competency or source of competitive advantage often need architectural control and IP protection, favoring framework-based approaches. On the other hand, organizations pursuing AI primarily for operational efficiency may be more willing to accept platform constraints in exchange for faster deployment.
- **Time-to-market expectations**
  Urgent timelines (i.e., weeks to two months) generally favor managed platforms for speed and reduced build effort. Hybrid approaches suit moderate timelines (three–six months), while longer investment horizons (greater than six months) enable custom builds that prioritize strategic differentiation over speed.

Technical

Technical evaluation determines whether a given solution can meet customization, integration, performance, and regulatory requirements.

Key considerations include:

- **Customization and architectural control**
  Frameworks offer full control over agent architecture design, orchestration, and workflows, making them ideal for complex or non-standard processes. Platforms trade this flexibility for ease of use, which can accelerate deployment but limit customization.
- **Integration complexity**
  Platforms typically offer pre-built connectors for mainstream enterprise systems (e.g., Salesforce, ServiceNow, Microsoft 365), reducing deployment timelines by up to 40%. Frameworks may be a better choice for legacy system or proprietary integrations.
- **Scalability and performance**
  Platforms provide built-in scalability and infrastructure management, making them suitable for variable workloads. Frameworks enable fine-tuned optimization that can outperform platforms in throughput and latency for high-volume or mission-critical workloads.
- **Security and compliance requirements**
  Platforms often include pre-built certifications and guardrails for standard compliance. Frameworks are necessary for organizations with unique regulatory requirements, data sovereignty constraints, or highly specialized security policies.

Financial

Financial evaluation must extend beyond initial implementation costs to a multi-year horizon, generally three years, that accounts for total cost of ownership.

Key considerations include:

- **Up-front and ongoing investment**
  Custom builds in framework-based approaches require significant engineering investment—often three times higher than platform-based solutions—along with ongoing maintenance and infrastructure costs. Platforms bundle infrastructure hosting, updates, and support into subscription or consumption-based pricing models.
- **Vendor dependency and pricing risk**
  Platform pricing models can simplify budgeting but introduce long-term vendor lock-in and exposure to pricing changes. Frameworks reduce dependency on a single vendor but shift cost and risk internally.
- **Return on investment and payback**
  Organizations should expect pilot initiatives to deliver 10–25% efficiency gains, with mature implementations scaling to 30–50% improvements as use cases expand and operating models stabilize.

Operational

Operational evaluation focuses on the organization's ability to support, monitor, and evolved agentic AI solutions over time.

Key considerations include:

- **Maintenance and support burden**
  Framework-based approaches place full responsibility for updates, scaling, and security on internal teams, requiring strong DevOps and MLOps capabilities. Platforms shift much of this burden to vendors through service-level agreements and automated updates.
- **Observability and debugging**
  Frameworks enable deep, customizable observability and tracing, which is critical for troubleshooting complex workflows and advanced optimization. Platforms provide built-in monitoring tools (e.g., AWS CloudWatch, Azure Monitor) that are sufficient for standard debugging but may limit deeper inspection.

Applying the criteria

The relative weighting of these four evaluation dimensions will vary based on organizational context. Factors such as market conditions, regulatory environment, competitive landscape, available talent and skills, and existing technology investments all influence how trade-offs should be assessed.

Once weightings are defined, organizations can score technology options against the criteria to develop a short list of viable candidates. These options should then be validated through targeted proof-of-concept pilots tied to clearly defined use cases and measurable outcomes. From there, organizations can move confidently from analysis to execution.

Conclusion

Now is the time to move from exploration to execution, as agentic AI moves from experimentation to enterprise expectation. As the market proliferates, the organizations that create durable value will not be those that adopt the most tools, but those that make disciplined technology choices aligned to how they operate.

After making an honest assessment of your current state and prioritizing your use cases, the next step is to design your future state and select the right agentic AI technology foundation. This decision will shape your organization's ability to scale, govern, and integrate agentic AI into core operations effectively.

Organizations should begin by shortlisting technology options and vendors that align with their strategic priorities, technical requirements, and compliance obligations, then validating those choices through targeted proof-of-concept pilots that demonstrate measurable business value. These pilots should be designed to test not only technical performance, but also operating-model fit and scalability.

In practice, translating selection decisions into results requires disciplined evaluation, operating-model alignment, and structured execution. At Alvarez and Marsal, our AI & Analytics team works with organizations across these dimensions by helping define and weigh criteria, assess trade-offs, structure pilots, and establish governance models that enable agentic AI to scale securely and sustainably.

Whether this process is led internally or supported by an experienced, unbiased partner, act decisively. The organizations that establish a robust and well-aligned agentic AI foundation today will be best positioned to scale, govern, and realize value as agentic AI adoption accelerates.

Contact Us

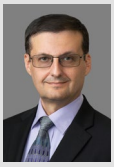

Dan Simion
Managing Director
dsimion@alvarezandmarsal.com

David Dina
Senior Director
ddina@alvarezandmarsal.com

Laura Gibbs
Senior Director
lgibbs@alvarezandmarsal.com

Follow A&M on:

ABOUT ALVAREZ & MARSAL

Founded in 1983, Alvarez & Marsal is a leading global professional services firm. Renowned for its leadership, action and results, Alvarez & Marsal provides advisory, business performance improvement and turnaround management services, delivering practical solutions to address clients' unique challenges. With a worldwide network of experienced operators, world-class consultants, former regulators and industry authorities, Alvarez & Marsal helps corporates, boards, private equity firms, law firms and government agencies drive transformation, mitigate risk and unlock value at every stage of growth.

To learn more, visit: AlvarezandMarsal.com

ALVAREZ & MARSAL
LEADERSHIP. ACTION. RESULTS.℠