BENCHMARK REPORT

2026

Agentic AI Readiness

Fern Halper, Ph.D.,
TDWI VP of Research



tdwi

Contents

TDWI Benchmark: Agentic AI Readiness

By Fern Halper, Ph.D.

Introduction: The Agentic Wave Is Coming 3

Agentic AI Adoption 4

What Companies May Underestimate About Agentic AI . . . 4

Agentic AI Readiness Framework 5

Results from the Agentic AI Readiness Assessment 7

Organizational and Capability Readiness 8

Data and Context Readiness. 11

Technology and Engineering Readiness 14

Governance, Control Readiness. 15

Operationalization Readiness 18

Are Organizations Ready? 19

Conclusion 20

About the Author. 21

Methodology 21

About the Sponsors 22

© 2026 by TDWI, a division of 1105 Media, Inc. All rights reserved. Reproductions in whole or part are prohibited except by written permission. Email requests or feedback to info@tdwi.org. Product and company names mentioned herein may be trademarks and/or registered trademarks of their respective companies. Inclusion of a vendor, product, or service in TDWI research does not constitute an endorsement by TDWI or its management. Sponsorship of a publication should not be construed as an endorsement of the sponsor organization or validation of its claims. This report is based on independent research and represents TDWI's findings; reader experience may differ. The information contained in this report was obtained from sources believed to be reliable at the time of publication. Features and specifications can and do change frequently; readers are encouraged to visit vendor websites for updated information. TDWI shall not be liable for any omissions or errors in the information in this report.

Introduction: The Agentic Wave Is Coming

Artificial intelligence is entering a new phase. What began as a rapid increase in generative AI adoption, often driven by accessible tools that can produce text, code, and content, has evolved into something more: systems that do not just generate, but act. In this next wave, often referred to as agentic AI, organizations are moving toward systems that can reason, plan, and execute tasks across enterprise environments.

The shift is significant. Generative AI lowered the barrier to entry for AI by enabling organizations to experiment quickly, often without deep investments in data or infrastructure. As a result, adoption surged. However, many of these early deployments remained limited—focused on productivity gains such as drafting content or summarizing information. Agentic AI, by contrast, introduces the possibility of automating entire workflows and decision processes. It represents a move to coordinated systems that operate across tools, data, and business processes. This promises potentially huge transformations and disruptions.

At its core, agentic AI builds on the same foundational technologies as generative AI but extends them with additional capabilities. Rather than responding to a prompt and producing an output, an agentic system operates through a continuous loop of perceiving, reasoning, acting, and learning. It gathers information from its environment, interprets that information in context, determines a course of action, executes against systems or tools, and evaluates the outcome.

In practice, this means that agentic systems can carry out end-to-end tasks. For example, instead of simply summarizing supplier contracts, an agentic system might retrieve those contracts, analyze pricing trends, compare them against external market data, and recommend or even initiate renegotiation workflows. These systems are not just advisory; they are operational.

A defining characteristic of agentic AI is the use of agents: modular components responsible for specific tasks and orchestration mechanisms that coordinate how those agents interact. In more advanced use cases, multiple agents collaborate within a structured workflow, each contributing specialized capabilities such as data retrieval, analysis, communication, or execution. This multi-agent paradigm enables organizations to break down sophisticated processes into coordinated steps, but it also requires careful design, governance, and control.

This is where agentic AI diverges most clearly from generative AI. Whereas generative AI systems primarily produce outputs, agentic systems take action. They interact with APIs, trigger workflows, update systems, and influence real-world outcomes. As a result, the stakes are higher and the risk surface often greater. Errors are no longer confined to incorrect answers; they can propagate through workflows, affect downstream decisions, and create operational risk.

The reality is that many organizations are still in the early stages of this journey. While the concept of agentic AI is gaining traction,

true enterprise-grade implementations (particularly multi-agent systems operating at scale) remain limited. Many organizations are experimenting with single-agent systems or extending generative AI use cases with limited automation. Fully orchestrated, autonomous systems are still emerging.

Agentic AI Adoption

Interest in agentic AI is high, but current adoption remains uneven. TDWI research indicates that while organizations are actively exploring agentic AI, relatively few have moved beyond experimentation into production environments. In fact, less than 10% of organizations report having multi-agent systems in production, underscoring the gap between ambition and execution.[1]

Organizations often begin with accessible tools, in this case, off-the-shelf generative AI applications, and then move toward more customized solutions that incorporate enterprise data. From there, they experiment with retrieval-augmented generation (RAG) and single-agent systems before attempting more complex, orchestrated workflows.

However, each step introduces additional complexity. Moving from generative AI to agentic AI is not simply an incremental change; it requires a different level of architectural maturity, governance, and operational discipline. As a result, many organizations remain in pilot or proof-of-concept stages. In TDWI research, while a significant portion of organizations report building custom AI assistants or experimenting with agent-based systems, far fewer have operationalized these capabilities at scale.

In many cases, AI adoption is a progression. Organizations are moving along a maturity curve, gaining experience with simpler systems before advancing to more complex ones. The challenge is that many underestimate the effort required to make this transition successfully.

What Companies May Underestimate About Agentic AI

In the qualitative research for this report, one theme that emerged was that organizations tend to overestimate their readiness for agentic AI and underestimate what it actually requires. This includes the data foundation, managing new data types, and the importance of context, skills, and governance. Across the quantitative results, the same pattern was evident: organizations are building capabilities, but that progress is uneven and not yet integrated into a cohesive, enterprise-ready foundation. In many cases, organizations have parts of the system in place, but not the system itself yet.

The data foundation is one significant gap. Many organizations assume that existing investments in analytics or machine learning are sufficient. In reality, agentic systems place fundamentally different demands on data. It is no longer enough for data to be available; it must be continuously reliable, governed, and enriched with context to support autonomous decision-making. For instance, in agentic systems, errors can compound. A small issue in the data quality in one step of a workflow can cascade into subsequent steps, leading to issues that can be harder to correct.

Organizations also underestimate the importance of context. Agentic systems

[1] See the Q1 2026 *TDWI Blueprint Report: Building Agentic and Generative AI: Enterprise Data Foundations and Applications,* available at https://tdwi.org/blueprint.

require not just data, but semantic understanding, i.e., understanding the business logic and intent. That includes definitions, relationships, and business meaning that enable accurate interpretation. Without this, even high-quality data can produce unreliable results.

Another area that is often overlooked is the complexity of working with unstructured data. While generative AI has made it easier to access and analyze text, organizations frequently lack the processes and metrics needed to ensure its quality and relevance. Issues such as chunking, context windows, and preprocessing become critical, yet are not well understood.

Skills and organizational readiness can also be major constraints. Many organizations lack the expertise needed to design, deploy, and manage agentic systems. This includes not only technical skills, but also the ability to interpret system behavior, debug workflows, and implement governance controls.

Governance itself is frequently underestimated. As systems become more autonomous, the need for oversight, policy enforcement, and risk management increases. Yet many organizations rely on ad hoc approaches or assume that existing controls are sufficient. In reality, agentic AI requires new governance models that account for dynamic, multi-step workflows and system-level interactions.

Finally, organizations often underestimate the operational challenge. Moving from a successful pilot to a production system is a significant leap. It requires integrating AI into real workflows, managing scale, ensuring reliability, and maintaining trust. Many initiatives stall at this stage, not because the technology fails, but because the supporting capabilities are not in place.

A key issue is that organizations often equate the ability to build or experiment with agents with readiness. In practice, agentic systems operate under very different conditions than traditional analytics or generative AI. They require continuous access to trusted, contextualized data, well-defined processes, and clear constraints on behavior. Without these elements, systems that appear effective in isolated scenarios may break down when exposed to real-world complexity. This contributes to a broader pattern seen in the data: organizations may perceive themselves as further along than their underlying readiness supports.

This report will dive into these areas in more detail.

Agentic AI Readiness Framework

How ready are organizations for agentic AI?

To answer this question, TDWI developed the Agentic AI Readiness Assessment, a framework designed to evaluate an organization's ability to move from experimentation to enterprise-scale deployment using five dimensions.

- **Organizational and Capability Readiness.** This dimension assesses whether the organization has the people, leadership alignment, and culture needed to support agentic AI. It includes factors such as executive sponsorship, strategic clarity, funding, availability of skilled talent, AI literacy across the organization, and the ability to understand and manage AI-driven

systems. It answers the question: Does the organization have the capabilities and alignment required to successfully adopt and manage agentic AI?

- **Data and Context Readiness.** This dimension evaluates whether the organization's data can effectively support agentic AI systems that reason and act with context. It includes data availability, quality, consistency, integration across systems, support for structured and unstructured data, and the presence of semantic meaning and metadata to provide business context. It answers the question: Is the organization's data trustworthy, accessible, and sufficiently contextualized to power agentic AI?
- **Technology and Engineering Readiness.** This dimension examines whether the technical environment/architecture are capable of supporting agentic AI and whether the engineering talent is in place to do so. It includes platforms and infrastructure, integration capabilities, support for orchestration and tool use, scalability, and the engineering resources required to build and maintain agentic systems. It answers the question: Can the organization's technology stack and engineering know-how support the development and execution of agentic AI workflows at scale?
- **Governance, Risk, and Control Readiness.** This dimension assesses the organization's ability to manage the risks associated with agentic AI and ensure responsible use. It includes policy frameworks, data access controls, monitoring and observability, auditability, compliance, and mechanisms to enforce guardrails on AI behavior. It answers the question: Does the organization have the controls and governance needed to safely deploy and manage agentic AI systems?
- **Operationalization and Learning Readiness.** This dimension evaluates the organization's ability to move agentic AI from experimentation into production and continuously improve it over time. It includes deployment processes, workflow integration, human-AI interaction models, monitoring and feedback loops, and the ability to scale and refine systems based on performance. It answers the question: Can the organization successfully operationalize, scale, and continuously improve agentic AI in real-world use cases?

These dimensions were designed this way because agentic AI is not just a technology challenge. It is an enterprise challenge. Success depends on the alignment of data, architecture, governance, skills, and organizational processes. Organizations may be advanced in one dimension and lagging in another. For example, a company may have strong leadership support but lack the data infrastructure needed to support agentic systems. Or it may have technical capabilities but insufficient governance. That can be a problem.

Agentic AI has the potential to fundamentally change how work gets done by automating processes and reducing cycle times. On the other hand, it raises critical questions about readiness. Do organizations have the data foundations, governance structures, skills, and operational capabilities required to support these systems? The assessment was designed

to explore these questions further and benchmark where organizations are today.

Results from the Agentic AI Readiness Assessment

Where are organizations now in terms of agentic AI readiness? The overall results of the TDWI Agentic AI Readiness Assessment indicate that organizations are making meaningful progress but are not yet fully prepared to operationalize agentic AI at scale. The median total score across respondents is 69 out of a possible 100, with each of the five dimensions contributing up to 20 points. This puts most organizations above the midpoint threshold that indicates some readiness, but still short of the level associated with guaranteed readiness. In terms of stages of agentic AI readiness, they are primarily preparing—meaning that organizations are actively defining and aligning capabilities (Figure 1). In some areas they have moved into the next stage in terms of putting core foundations in place. However, there is a bit of a twist in this case in that organizations are still working on core components such as data and governance foundations.

Looking across the five dimensions, readiness is uneven (Figure 2). This unevenness reflects how organizations are approaching agentic AI. Capabilities such as technology infrastructure and operational processes are advancing more quickly, while foundational elements such as data consistency, governance enforcement, and organizational alignment

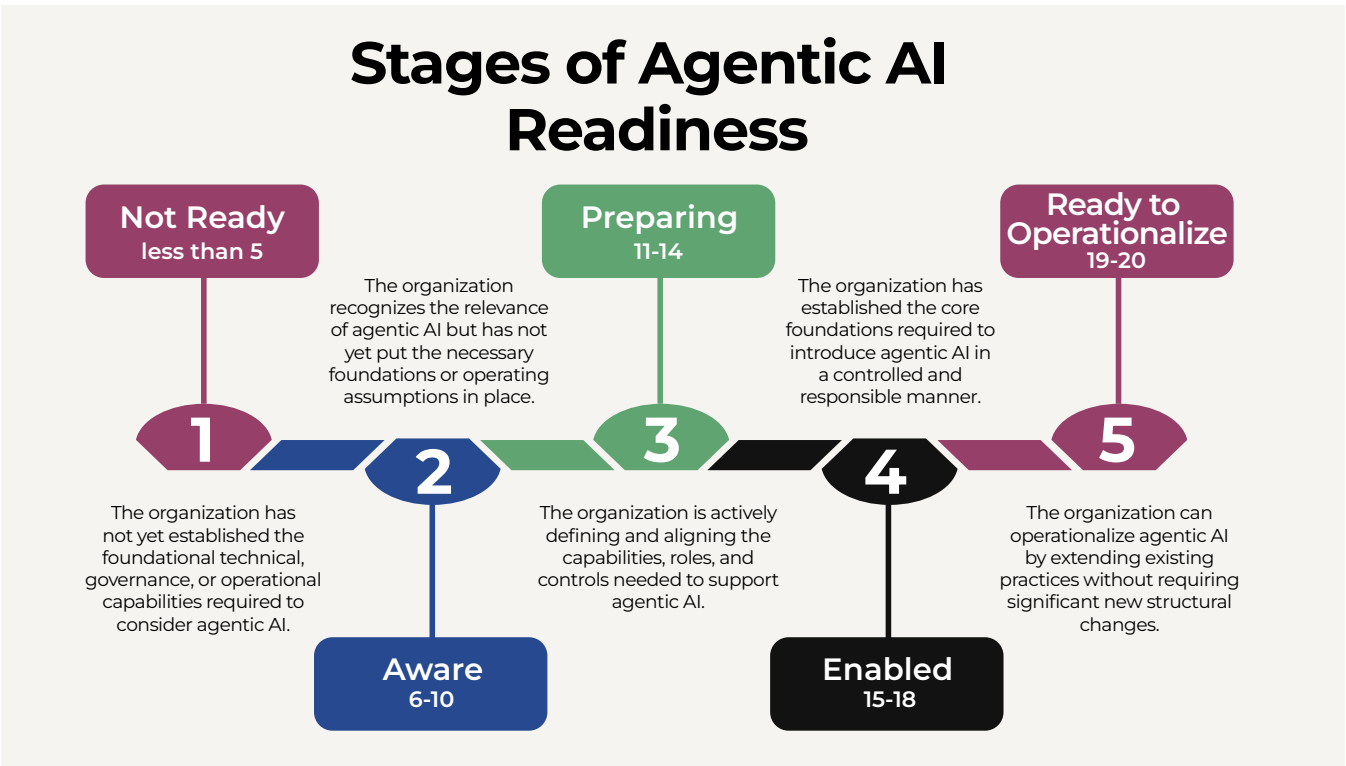


Figure 1. *Stages of readiness in TDWI's Agentic AI Readiness Assessment and the score range for an individual dimension in each stage.*

lag behind. Because agentic AI depends on the interaction of these components, gaps in any one area can limit the effectiveness of the whole. Taken together, the results show that while organizations are building momentum, they must continue to strengthen foundational areas, particularly data, governance, and skills, to fully realize the potential of agentic AI.

Let's look at each dimension individually.

Area	Median Score
Organizational and Capability Readiness	13
Data and Context Readiness	13
Technology and Engineering Readiness	15
Governance, Risk, and Context Readiness	14
Operationalization and Learning Readiness	15

Figure 2. *Median scores for each dimension of TDWI's Agentic AI Readiness Assessment based on 161 respondents. Overall median score is 69.*

Organizational and Capability Readiness

With a median score of 13 out of 20, organizational readiness is still in the preparation stage of readiness. This means that many respondents are defining capabilities, while others are putting core organizational capabilities in place. There is growing alignment around AI's importance, initial use cases are being identified, and some structural elements are emerging. The results from this assessment indicate that organizational readiness is uneven. Many organizations have not yet established the specialized skills and operational discipline required to support agentic AI at scale. In other words, the intent is clear, but readiness is still evolving.

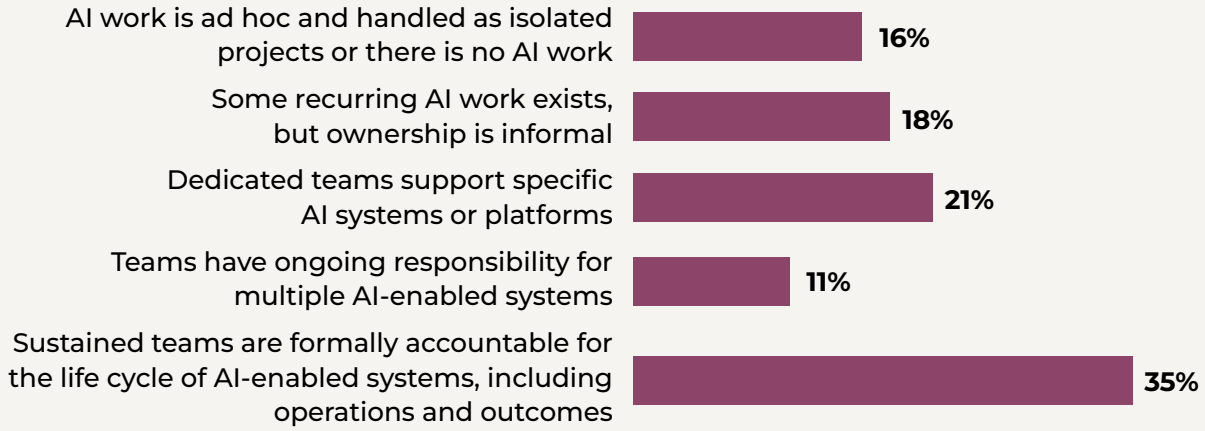
Areas of strength include that organizations are beginning to put foundational elements, particularly around establishing teams and funding, in place. In terms of teams, over one-third of respondents report having formally accountable teams responsible for the life cycle of AI systems, and another 21% have dedicated teams supporting specific systems or platforms (Figure 3). For funding, about a third report sustained funding across experimentation, deployment, and operations, suggesting that AI is moving beyond isolated pilots into more durable initiatives. Another 22% report that budget exists for both experimentation and production. These are good indicators that organizations are starting to view agentic AI as an ongoing capability rather than a one-off project.

Despite progress in structure and funding, organizations are less mature in the capabilities required to execute agentic AI effectively (Figure 4). Slightly less than one-third report that agentic AI skills are formally in place and supported strategically, while the majority rely on fragmented expertise across individuals, teams, or partners.

Process maturity is similarly uneven: just 27% say their target processes are comprehensively documented and understood, while the rest operate with incomplete, outdated, or poorly defined processes. This creates a risk: agentic AI systems depend on well-

Organizational and Capability Readiness Areas of Strength

In my company, there are teams with sustained responsibility for AI-enabled systems, not just short-term projects.



In my company, budget exists for both experimentation and delivery of agentic AI applications.

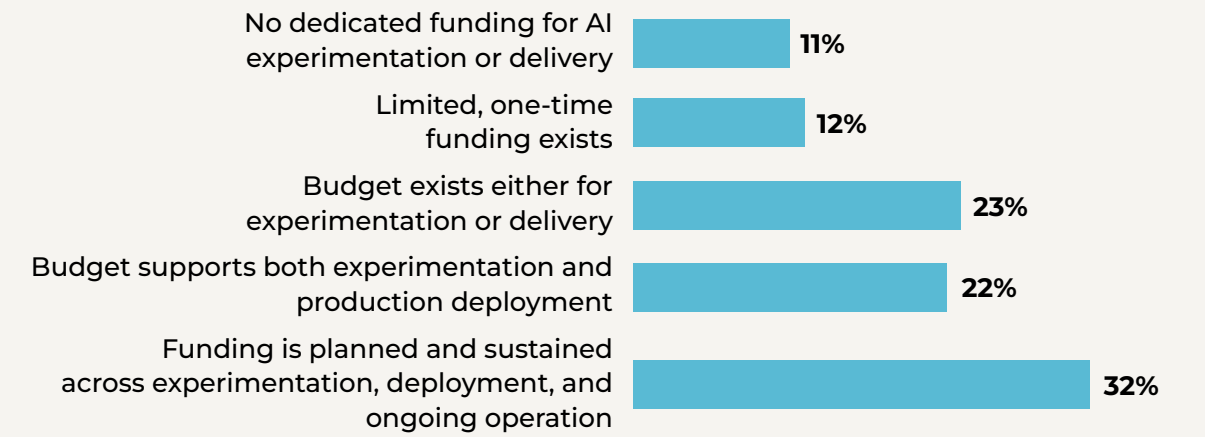


Figure 3. *Based on 161 respondents.*

understood processes, and organizations that lack this clarity may automate inefficiencies or introduce unintended consequences.

A final gap lies in organizational change readiness. While about a third of respondents report formal change management processes, most organizations are still early: 26% have only discussed impacts informally, and 9% have not considered them at all (not shown). This is an important issue for agentic AI. Unlike earlier analytics approaches, agentic systems alter workflows, decision-making authority, and human roles. Without deliberate change management, even technically successful

Organizational and Capability Readiness Areas for Improvement

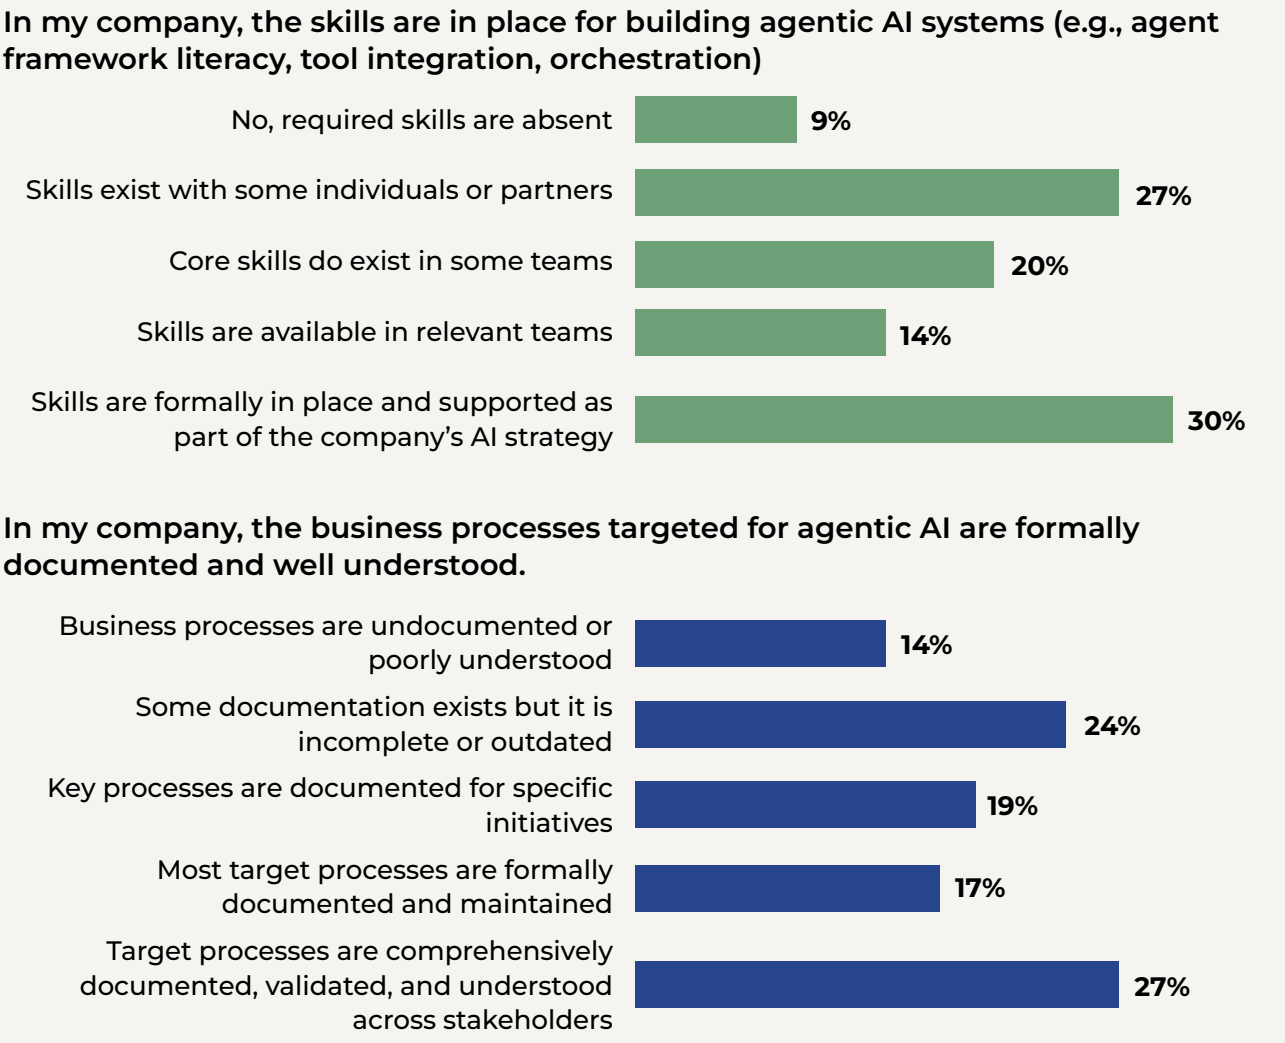


Figure 4. *Based on 161 respondents.*

implementations may struggle to gain adoption or deliver value.

Organizational readiness, then, is advancing but incomplete. Companies are beginning to invest, organize, and align around AI, but they have not yet built the organizational backbone required for agentic systems. The largest gaps in skills, process clarity, and change management are the areas that will matter for scaling autonomous, multi-agent workflows. These gaps are not just about capacity, but about depth of understanding. Agentic systems require teams that can interpret system behavior, diagnose failures, and understand how decisions are made across workflows. Without this, organizations may be able to experiment, but they will struggle to operationalize and scale these systems.

Data and Context Readiness

With a median score of 13 out of 20, data readiness also falls in the Preparing stage, behind the three remaining dimensions. Respondents have made progress in modernizing their data environments and recognize the importance of data for AI. This is not surprising, as in TDWI research we see that putting the data foundation in place for AI remains the top priority for organizations. We will see in the next section that the technology stack for AI is an area of strength for organizations.

> **Data readiness is key to agentic AI. Organizations are investing in metadata, access controls, and lineage, but these are not yet consistently integrated into a unified, trusted data foundation.**

However, for many, while areas of readiness exist around metadata, access controls, and lineage, these capabilities are not yet uniformly applied. This lack of consistency is an important issue for agentic AI, which depends on trusted, well-governed, and context-rich data.

On the positive side, 29% of respondents report comprehensive, up-to-date metadata for AI-relevant assets, another 19% say metadata is captured and governed, and 16% claim metadata is available for key data assets and use cases (Figure 5). Similarly, 50% report broadly implemented or fully automated fine-grained access controls, and 32% indicate comprehensive, automated lineage with 16% stating it is available for critical data flows and key systems. These are important building blocks for agentic AI, where systems must operate with clear data context, enforce permissions, and support traceability.

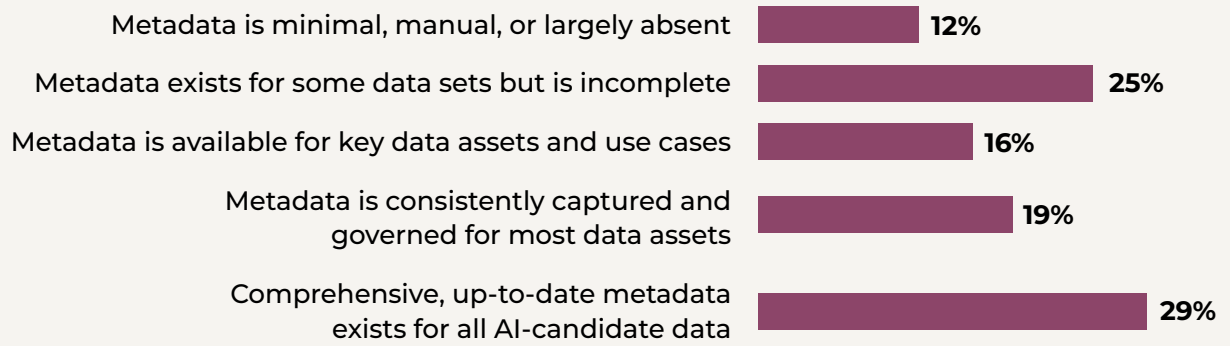
Despite progress, confidence in data, especially across different types, remains uneven (Figure 6). For structured data, only 47% report broadly trusted or enterprise-authoritative data, while the rest cite integrity issues, fragmentation, or limited coverage. The picture is the same or slightly more mixed for unstructured data. This suggests that many organizations are still struggling to manage and govern the diverse data types required for agentic AI, particularly for text, documents, and other unstructured sources that agents rely on heavily.

Another critical gap is the lack of a fully adopted semantic layer. Only 27% report a governed, enterprise-wide semantic layer that is machine-consumable, while the rest operate with partial, domain-specific, or inconsistent definitions (not shown). This is a significant limitation for agentic AI. Agents rely on consistent business meaning to reason, act, and coordinate across systems. Without a shared semantic foundation, organizations risk inconsistent outputs, misinterpretation of data, and reduced trust in AI-driven decisions.

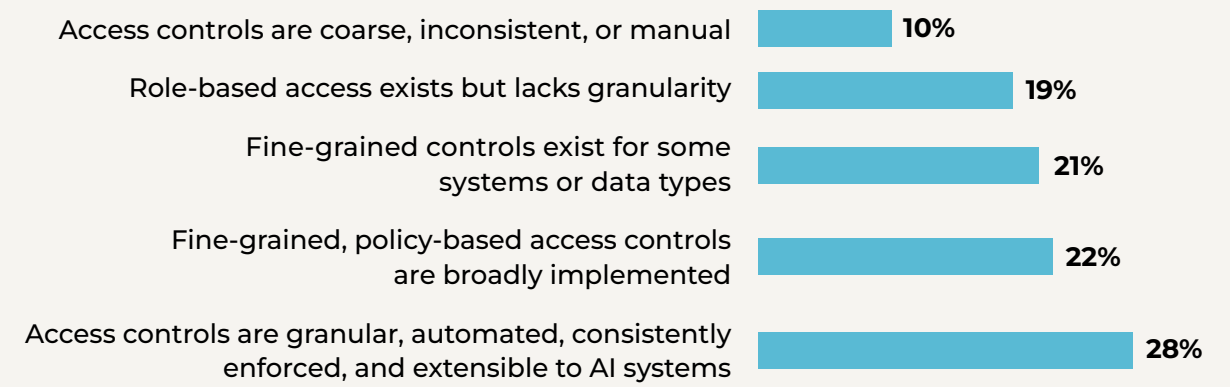
Overall, data readiness is key to agentic AI, and it is advancing. Organizations are investing in key technical capabilities such as metadata, access controls, and lineage, but these are not consistently integrated into a unified, trusted data foundation. The biggest gaps include data consistency, unstructured data

Data and Context Readiness Areas of Strength

In my company, metadata exists for data assets that would be candidates for use by AI systems, including agents.



In my company, fine-grained access controls are in place and could be applied to AI systems, including agents, for controlled read and write access to data.



In my company, decisions made by AI-driven systems could be traced back to the data used, based on existing lineage and traceability capabilities.

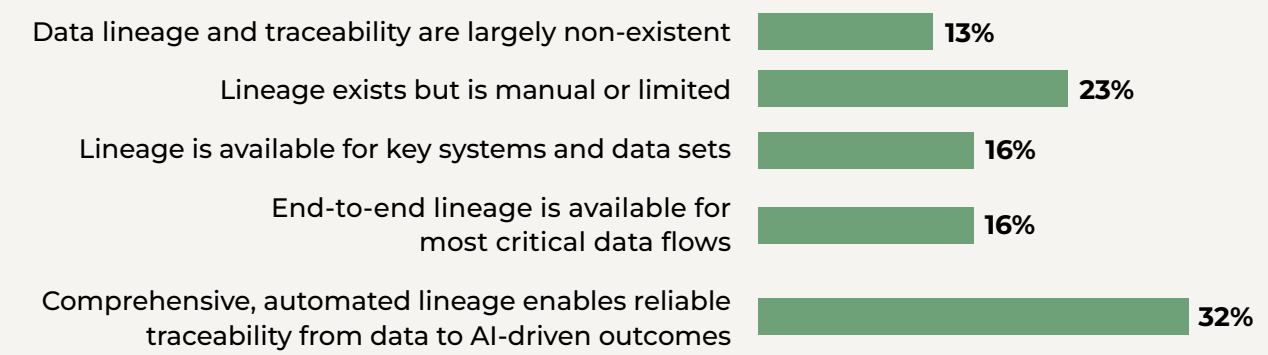


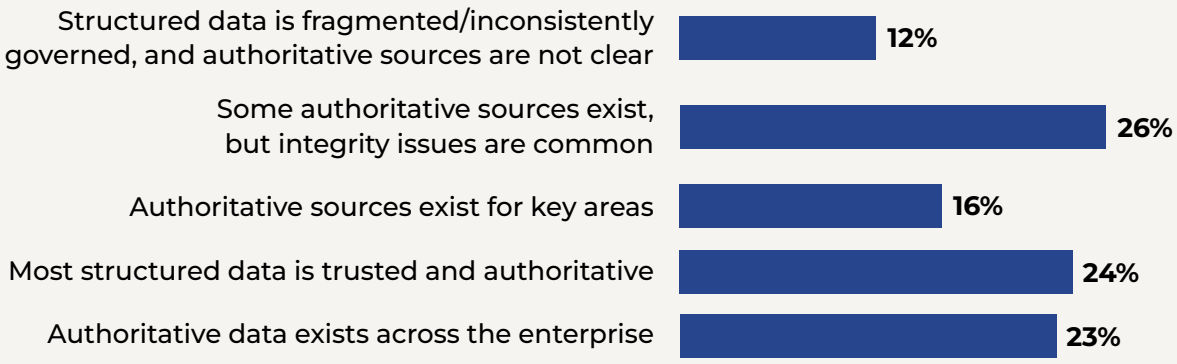
Figure 5. *Based on 161 respondents.*

governance, and semantic alignment, and these are some of the areas that agentic AI will depend on most. As a result, many organizations are at risk of building agentic systems on data foundations that are not yet robust enough to support them.

Additionally, in agentic environments, data issues do not occur in isolation. They compound across multi-step workflows. Small inconsistencies, missing context, or gaps in governance can propagate through systems, leading to increasingly unreliable outcomes. This makes data readiness not just a prerequisite, but an ongoing operational discipline that must be maintained continuously.

Confidence in Data

In my company, we are confident in the integrity of our structured data and have authoritative data sources that could be used by AI systems, including agents.



In my company, we are confident in the integrity of our unstructured data sources and have authoritative sources that could be used by AI systems, including agents.

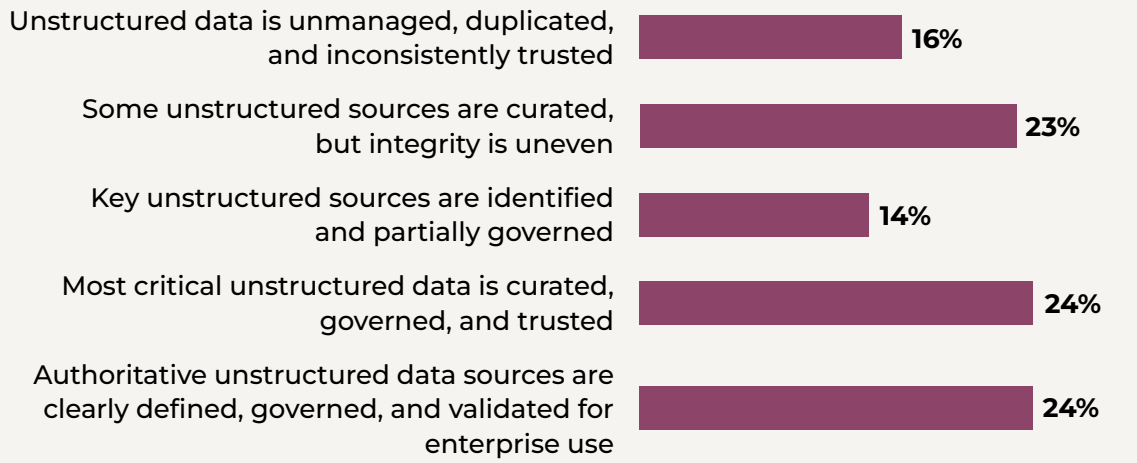


Figure 6. *Based on 161 respondents.*

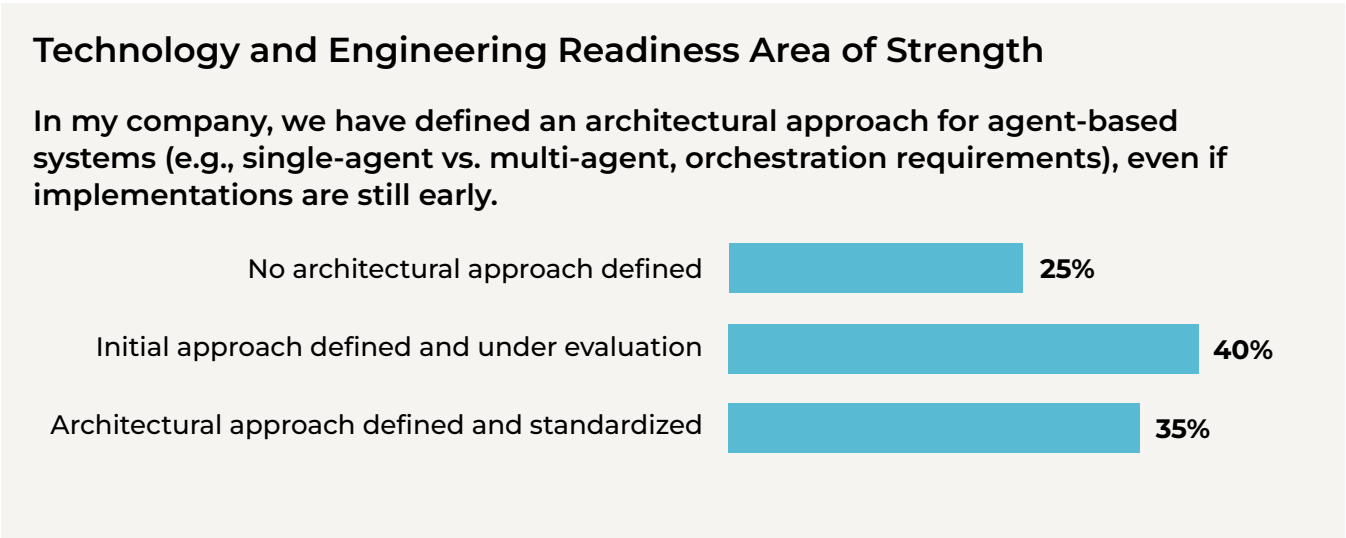


Figure 7. *Based on 161 respondents.*

Technology and Engineering Readiness

Technology and engineering readiness is one of the higher-scoring dimensions, with a median score of 15 out of 20, reflecting investment in infrastructure, tooling, and experimentation. This places it in the Enabling stage where organizations are putting capabilities in place to ultimately deploy agentic AI in a responsible way.

Organizations are building the technical foundations required for agentic AI, with many reporting modern cloud environments and early architectural thinking. This is not surprising, as we've seen this in previous TDWI research. However, while technical components are being assembled, they are not yet consistently integrated, standardized, or production-hardened. In effect, many organizations are technically capable of building agentic systems, but not yet able to build or operate them reliably at scale.

Organizations have made meaningful progress in standing up the technical stack for agentic AI. A majority (58%) report that cloud and compute infrastructure is provisioned and architected to support continuous, low-latency execution of agentic AI workloads (not shown). There is also evidence of growing architectural maturity (Figure 7).

Approximately 35% of respondents have an architectural approach defined and standardized; 40% claim to have an initial architectural approach defined or under evaluation. This indicates a shift toward more deliberate system design. More than half have evaluated or selected agentic frameworks (not shown).

While some of the infrastructure is in place to support agentic AI, there are weaknesses. For instance, integration remains a significant constraint. Only 22% of respondents report having assessed integration patterns and dependencies for compatibility with agent-based systems while the majority are still developing one or acknowledge the need without having acted. Even where integration exists, it is not fully reliable: just 45.3% consider

Confidence in Talent

In my company, we have access to engineering talent with hands-on experience in generative and agentic AI (internally or through partners).

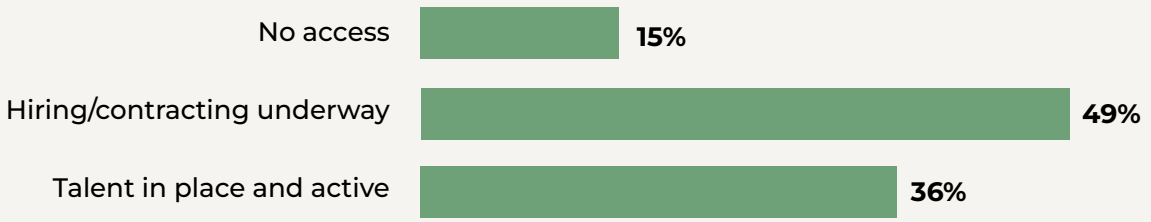


Figure 8. *Based on 161 respondents.*

it agent-ready, while nearly as many say it may not work effectively for agents (not shown). This is a critical issue, as agentic AI depends on reliable, real-time interaction with enterprise systems. Without this, even well-designed agents cannot operate effectively in production.

Finally, while many organizations report familiarity with agentic frameworks, deep engineering maturity is less widespread. For instance, while 36% of respondents believe they have the talent for agentic AI in place, close to 50% of respondents report actively hiring/contracting (Figure 8). This reflects a strong acknowledgement that organizations lack skills related to generative and agentic AI (discussed under Organizational Readiness), but need to put them in place.

A meaningful portion of respondents express uncertainty about their teams' ability to evaluate and implement these technologies, and only about half report a strong understanding of failure modes and mitigation strategies for agentic systems (not shown).

This points to a gap between tool adoption and engineering rigor. Agentic systems require robust testing, observability, and resilience—capabilities that are still developing in many organizations. It appears that organizations are moving toward putting the technology and ultimately the skills in place to build agentic systems. The challenge is whether they can integrate, manage, and sustain them in real-world environments.

Advances in tooling are making it easier to build agentic systems, including interfaces that allow users to define agents using natural language. However, this ease of development may create a false sense of readiness. While building agents is becoming simpler, ensuring they operate reliably, integrate with enterprise systems, and act on trusted data remains more complex.

Governance, Control Readiness

We've already discussed some governance readiness strengths. Governance and control readiness is progressing, with a median score of 14 out of 20, placing organizations

Governance and Control Readiness Areas of Strength

In my organization, policies and rules have been defined to govern agent behaviors and acceptable actions and these are accepted by legal and compliance teams.



In my company, real-time monitoring of AI and agent actions is in place.

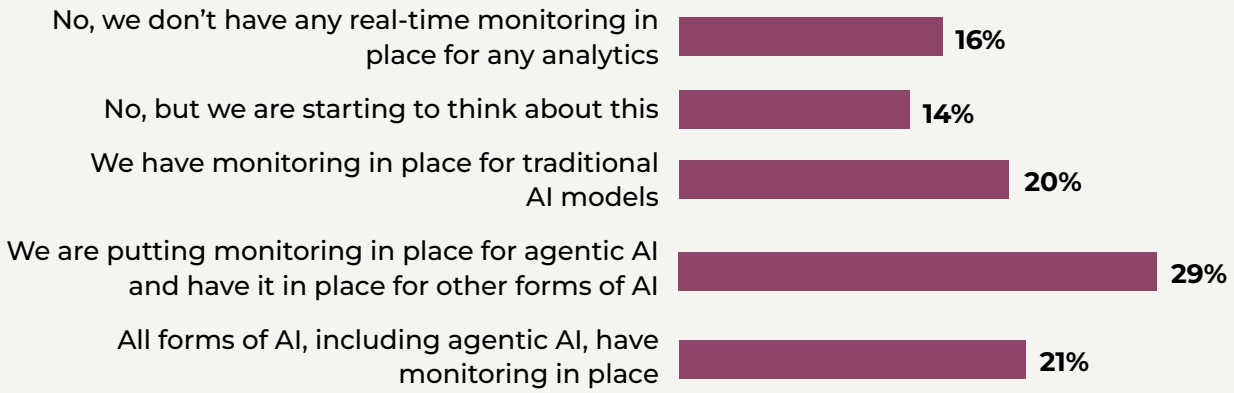


Figure 9. *Based on 161 respondents.*

at the tail end of the Preparing stage. Many organizations have begun to define policies, oversight models, and accountability structures for AI, including agentic systems. However, these efforts are often in progress rather than fully operationalized. Governance frameworks exist on paper or in pilot form, but are not yet consistently enforced across systems. This creates a gap: agentic AI introduces new levels of autonomy and risk, and without mature governance, organizations may struggle to maintain control, accountability, and trust.

Organizations are making clear progress in defining governance structures. For example, 42% report that policies governing agent behavior are fully approved, with another 37% actively drafting or piloting them (Figure 9). Similarly, over one-third (35%) have defined and approved decision boundaries for agents, with an additional 38% in progress (not shown). These results indicate that organizations are increasingly aware of the need for formal governance as AI systems become more autonomous. We have seen this in other TDWI research. Some organizations are putting the brakes

Governance and Control Readiness Area for Improvement

In my organization, clear ownership and accountability are assigned for agent-based systems.

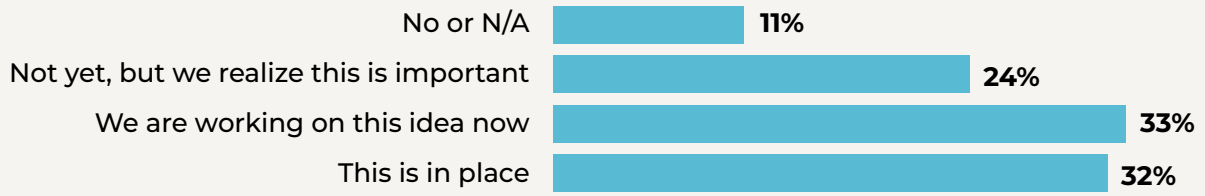


Figure 10. *Based on 161 respondents.*

on system deployments until they have the guardrails in place to ensure that agentic systems are being well governed.

There is also evidence that organizations are beginning to build operational governance capabilities via real-time monitoring and observability. This will be critical for agentic AI. Over 50% of respondents report real-time monitoring is already in place or they are putting it in place for agentic or other AI systems. As mentioned above, lineage is in place and about one-third of organizations surveyed are putting traceability fully in place with about the same amount actively claiming to be working on it (not shown).

Human oversight models are also evolving: about 30% have defined oversight across all AI types, while others are extending governance from generative AI to agentic systems (not shown). These are essential capabilities for managing risk and ensuring accountability in autonomous systems.

Despite progress, many governance capabilities remain incomplete or inconsistent. Only 32% report clear ownership and accountability in place, while the rest are still defining or recognizing the need for it (Figure 10). Only about a quarter report that autonomy levels are fully defined and approved, while the majority are still developing or have not yet addressed this. Similarly, oversight models are uneven, with many organizations either applying them only to generative AI or still in early stages of development (all not shown). This is a significant concern, as agentic AI systems require clear boundaries and escalation paths to operate safely.

Similarly, mechanisms for controlling agent behavior, such as the ability to pause or override actions (sometimes called the kill switch) are largely immature. This suggests that while organizations understand the need for governance, they have not yet translated this into enforceable, system-level controls, but perhaps they are ready for it.

The bottom line is that governance and control readiness is moving in the right direction for agentic AI. Organizations are actively defining policies, monitoring systems, and oversight models, but these capabilities are often incomplete or inconsistently applied.

Operationalization Readiness Area of Strength

In my company, existing AI or analytics operating models could be extended to support agent-based systems.

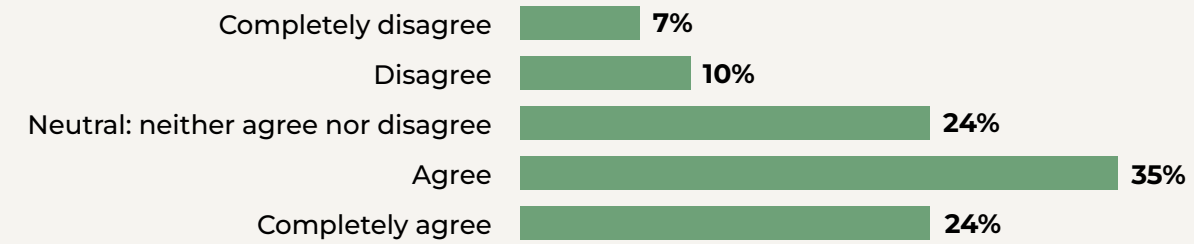


Figure 11. *Based on 161 respondents.*

Critical gaps, ownership, control mechanisms, and autonomy boundaries are the areas required to safely deploy agentic AI. Effective governance in agentic systems goes beyond defining policies; it requires translating those policies into enforceable, system-level controls. This includes the ability to monitor behavior in real time, intervene when necessary, and ensure that agent actions are constrained within clearly defined boundaries. Many organizations are still in the early stages of making this shift.

Operationalization Readiness

Operationalization readiness was the other highest-scoring dimension, with a median score of 15 out of 20, placing it in the Enabling stage. This suggests that many organizations are beginning to think seriously about how to extend existing AI capabilities into production environments for agentic systems. However, many are still in transition. While operating models, teams, and processes are being adapted, they are not yet consistently defined, formalized, or scaled. Organizations are thinking the right way, but not yet executing in production.

A strength is that organizations believe they can build on what they already have. Sixty percent agree or strongly agree that existing AI operating models can be extended to agent-based systems (Figure 11), and more than three-quarters say their engineering and data teams can be upskilled or extended to support these systems (not shown). Slightly over half of the respondents believe existing expectations around transparency, explainability, and oversight can be applied to agents. This indicates that organizations are not starting from scratch; they see agentic AI as an evolution of their current AI capabilities rather than a completely new paradigm. That is a good thing, although agentic AI will require more work.

Despite this progress, the dominant pattern across nearly all operational questions is that capabilities are still being thought through rather than fully implemented. This is the case for roles and responsibilities, performance metrics and oversight

processes, human oversight, and incident management (not shown).

Another important gap is the understanding of system dependencies and interactions. Only about half agree or strongly agree that dependencies between AI systems are well understood, while a large portion remain neutral or disagree. This lack of clarity is a significant risk for agentic AI, where systems interact dynamically across workflows and environments.

Without a strong understanding of dependencies, organizations may struggle to manage change, prevent cascading failures, or ensure reliable outcomes. The result is a gap between capability and execution: many have not yet established the discipline required to run agentic systems reliably at scale. A key challenge is that agentic systems depend on the coordination of multiple components: data pipelines, orchestration layers, tools, and governance mechanisms. When these elements are not designed to work together, systems may function in limited scenarios but break down under real-world conditions.

Are Organizations Ready?

So, are organizations ready for agentic AI? The move toward agentic AI is certainly happening fast. At TDWI, we see that it is a top investment area in 2026. The overall results of the assessment appear to show that organizations are making progress toward this goal.

However, the term agentic AI means different things to different people. To some, it is generative AI and creating an "agent," to others, it is multi-agent orchestrated systems. The data from this assessment presents a mixed view. While most organizations aren't in the place to deploy and operationalize agentic AI, they are getting ready.

Taken together, the results suggest that organizations are not unprepared, but they are also not fully ready. They are assembling the components required for agentic AI, but those components are often developed independently rather than as part of an integrated system.



Agentic systems depend on the coordination of multiple components. If these elements don't work together, systems may function in limited scenarios but break down under real-world conditions.

This creates an important dynamic. Organizations may appear advanced because they have capabilities in place, modern infrastructure, emerging operating models, or initial use cases. However, these capabilities are not yet consistently aligned. Agentic AI requires coordination across data, governance, technology, and organizational processes. Weakness in any one area can limit the effectiveness of the entire system.

The data also suggests that organizations may be further along in perception than in practice. While many report agentic AI activity, fewer have the underlying skills, data foundations, and governance structures required to support these systems at scale. This gap

between perceived and actual readiness will be important to monitor as adoption evolves.

Ultimately, readiness for agentic AI should not be viewed as a single score, but as the alignment of capabilities across the enterprise. The organizations that succeed will be those that move beyond isolated progress and focus on building a cohesive, scalable foundation—one where data, governance, technology, and operational processes function together as a unified system.

Conclusion

Organizations are clearly making progress toward agentic AI, but that progress is uneven and not yet fully integrated. Across the assessment, most respondents fall into a "preparing" stage. They are moving beyond early experimentation and beginning to establish key capabilities. Strengths are emerging in areas such as technology infrastructure and operational thinking, where organizations are leveraging existing investments and extending prior AI initiatives.

At the same time, foundational elements including data consistency, governance enforcement, and organizational capability remain less mature. This imbalance creates a situation where organizations may appear more advanced than they are, having assembled parts of the system without yet aligning them into a cohesive whole.

The results suggest that readiness for agentic AI is not defined by progress in any single dimension, but by the degree to which capabilities are coordinated across the enterprise. Organizations with higher levels of readiness tend to exhibit more consistent alignment across data, governance, technology, and operational processes, rather than isolated strengths in one area.

For most, the challenge is not initiating agentic AI efforts, but bringing these elements together in a way that supports reliable, scalable systems. As the market continues to evolve, the gap between perceived maturity and actual readiness will be a defining factor in how successfully organizations move from experimentation to production.

About the Author



FERN HALPER, PH.D., is VP of TDWI Research. Her work focuses on AI, generative AI, agentic AI, AI governance, cloud computing, and other modern analytics approaches. She has more than 25 years of experience in data and business analysis and AI and has published numerous articles on data mining and information technology. Halper is co-author of "Dummies" books on cloud computing, hybrid cloud, and service management, as well as *Big Data for Dummies*. Halper is also the author of the 2026 book, *Data Makes the World Go 'Round: The data, tech, and trust behind AI success*. She has been a partner at industry analyst firm Hurwitz & Associates and a lead analyst for Bell Labs. She has taught at both Colgate University and Bentley University.

Her Ph.D. is from Texas A&M University. You can reach her at fhalper@tdwi.org and on LinkedIn at https://www.linkedin.com/in/fbhalper/.

Methodology

The TDWI Agentic AI Readiness Assessment was fielded in early 2026 to help organizations evaluate how prepared they are to adopt and operationalize agentic AI and provide TDWI with a comprehensive view of current capabilities. The assessment was designed to measure readiness across five key dimensions: Organizational and Capability Readiness; Data and Context Readiness; Technology and Engineering Readiness; Governance, Risk, and Control Readiness; and Operationalization and Learning Readiness. Each dimension was scored on a 0–20 scale, with a total possible readiness score of 100.

A total of 161 respondents met the quality criteria for inclusion in this analysis. Responses were screened for completeness, consistency, and response quality, with incomplete or low-quality submissions removed to ensure the reliability of the results. Qualitative SME interviews were also conducted.

The respondent base represents a cross-section of industries, company sizes, and geographies. The largest industry segments represented include financial services, information technology, healthcare, and manufacturing, providing a balanced view across both data-intensive and operationally complex sectors. In terms of company size, the sample skews toward mid-sized and large enterprises, with the highest representation from organizations with annual revenues between $500 million and $999 million, followed by those in the $1 billion to $4.9 billion range and $10 billion or greater.

Geographically, the respondent base is predominantly located in the United States, with additional representation from Latin America, Europe, Asia, and Canada. While the sample includes global perspectives, the findings are most reflective of organizations operating in North America.

precisely

As a global data integrity provider, Precisely ensures that your data is accurate, consistent, and contextual. Our portfolio, featuring the Precisely Data Integrity Suite, brings together software, data, and data strategy consulting services. This unique combination enables organizations to move to Agentic-Ready Data, the highest quality of data that is integrated, governed, and enriched for AI, automation, and analytics initiatives at enterprise scale. Over 12,000 organizations in more than 100 countries, including 95 of the Fortune 100, trust Precisely to support some of the world's most complex, regulated, and mission-critical data environments. Learn more at precisely.com.

snowflake

Snowflake is the platform for the AI era, making it easy for enterprises to innovate faster and get more value from data. More than 12,000 customers around the globe, including hundreds of the world's largest companies, use Snowflake's AI Data Cloud to build, use, and share data, applications, and AI. With Snowflake, data and AI are transformative for everyone. Learn more at snowflake.com (NYSE: SNOW).

zoominfo

ZoomInfo (NASDAQ: GTM) is the trusted go-to-market platform that helps businesses find, acquire and grow their customers. It delivers real-time world-class data, insights, and technology to more than 35,000 companies worldwide. Businesses use ZoomInfo to increase efficiency, consolidate their technology stacks, and align sales and marketing teams in one platform. ZoomInfo is a leader in data privacy and emphasizes GDPR and CCPA compliance. In addition to creating the industry's first proactive notice program, it has earned numerous data security and privacy certifications. For more information about how ZoomInfo can help businesses grow their revenue at scale, please visit zoominfo.com.



research

TDWI Research provides research and advice for data professionals worldwide. TDWI Research focuses exclusively on data management and analytics issues and teams up with industry thought leaders and practitioners to deliver both broad and deep understanding of the business and technical challenges surrounding the deployment and use of data management and analytics solutions. TDWI Research offers in-depth research reports, commentary, inquiry services, and topical conferences as well as strategic planning services to user and vendor organizations.



Transforming Data With Intelligence™

A Division of 1105 Media
6300 Canoga Avenue, Suite 1150
Woodland Hills, CA 91367

E info@tdwi.org

tdwi.org