Fivetran | Redpoint INSIGHTS

The 2026 agentic AI readiness index

60% of enterprises are investing millions in agentic AI, but only 15% have the data foundation to run it securely and effectively at scale

Table of contents

Introduction

Agentic AI marks a clear break from earlier generations of artificial intelligence. Instead of simply predicting outcomes or generating content, it is designed to plan work, coordinate tasks, and take action across real business workflows.

But that shift changes the risk profile. When AI systems move from advising humans to acting on their behalf, the quality of the data underneath them matters far more. Agents rely on data that is current, governed, and trustworthy to make decisions in context.

Being "ready" for agentic AI in production means having 4 foundational conditions in place: dependable real-time data, transparent lineage, interoperability across systems, and governance controls that enforce security and sovereignty requirements. When data is incomplete, outdated, or poorly controlled, agentic AI makes mistakes faster and at greater scale.

Most enterprise data environments were never built for this level of autonomy. As companies push agentic AI into production, familiar problems like brittle pipelines, missing lineage, governance gaps, and uneven data quality move from background issues to becoming points of failure.

Gartner estimates more than 40% of agentic AI projects will be canceled by 2027, often because organizations move too quickly into use cases their data and infrastructure cannot reliably support.

Based on a global survey of 400 data professionals, this report explores where enterprises are investing, why their data foundations are breaking down, and what separates organizations that turn agentic AI into lasting value from those left with expensive experiments.

Executive summary

Agentic AI is drawing serious investment, but there's a growing gap between how fast organizations are moving and how ready their data foundations are to support the success of agentic AI. The following key findings show where risk is concentrated and why data foundations have become the primary bottleneck.

1. **High investment, limited readiness**
   Nearly 60% of organizations report investing millions to tens of millions in agentic AI, yet only 15% say they are totally prepared to support it with their current data foundation.

2. **Data problems pose the greatest barrier to achieving AI goals**
   Data quality and lineage issues (42%), sovereignty and regulatory compliance (39%), and security and privacy risk (39%) are the most frequently cited barriers to achieving agentic AI goals.

3. **Lack of interoperability drives vendor lock-in and** limits the flexibility needed to deliver agentic AI
   86% of data leaders say platform extensibility and interoperability are important or critical. Among those worried about lock-in, data integration is viewed as the single biggest lock-in risk, ahead of both cloud providers and major application vendors.

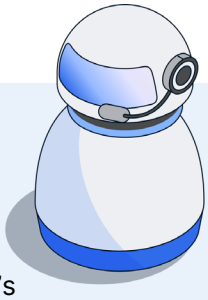

Agentic AI readiness index

"Agentic AI readiness index" is a composite score (0–100) that reflects how prepared an organization's data foundation is to support agentic AI in production. The index aggregates responses across 6 dimensions of AI readiness, which fall into core themes like data infrastructure, governance and security, interoperability, and operational readiness.

Scores are derived from how respondents rated their organization's readiness across these core pillars, then normalized into a single index to enable comparison across industries, company sizes, and adoption stages. Organizations were subsequently grouped into readiness segments (not prepared, somewhat prepared, fully prepared) based on their overall index score.

❖ KEY TAKEAWAY 1

Agentic AI investment is outpacing data readiness

Enterprises have already invested millions in AI initiatives, but the foundational data infrastructure required to power them is still lagging. As a result, agentic AI adoption has become an immediate business risk, limiting organizations' ability to realize meaningful ROI and achieve the outcomes they expect from AI investments.

Only **15%** of organizations say they are fully prepared for agentic AI today. Yet, most teams are moving forward while still working through gaps in the data foundation needed for secure, reliable production.

Readiness by segment

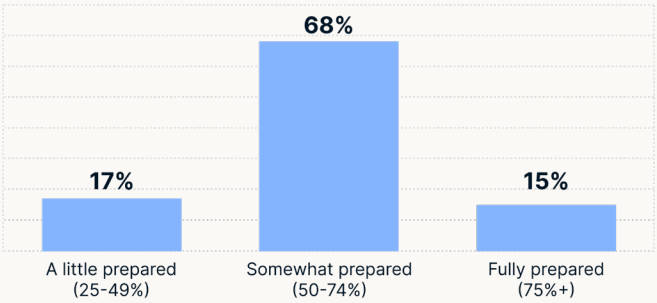

Average readiness index: 61%

(Based on Agentic AI readiness index)

The adoption stage distribution shows how widespread this readiness gap has become:

- **41%** of organizations are already using agentic AI in production, split between limited and scaled production
- **58%** are still evaluating or piloting agentic AI
- Only **1%** of organizations are not evaluating agentic AI at all

Stage of agentic AI use

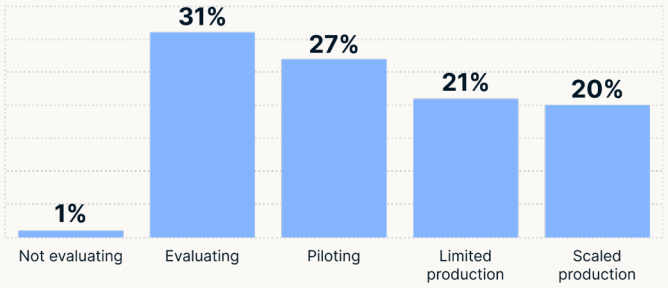

Production adoption is moving faster than data foundations can support. In many organizations, agentic AI reaches production before data reliability, lineage, and governance are in place to manage autonomous workflows safely, which is what being fully prepared requires.

60% of organizations say they have already invested millions to tens of millions in agentic AI.

Agentic AI is no longer confined to pilots, as many organizations are funding it at enterprise scale. Even organizations in scaled production report only **62%** readiness across the 4 data foundation pillars:

- Freshness and reliability
- Transparency and lineage
- Interoperability
- Sovereignty and governance

These 4 pillars underline the minimum conditions for agentic AI in production:

- Current and dependable data
- Lineage that shows where it came from and how it changed
- Interoperability across tools and clouds without brittle, custom work
- Governance controls that enforce sovereignty requirements by region

Weakness in any pillar limits what agents can safely automate and increases operational and compliance risk.

Agentic AI adoption by industry

	Technology	Finance	Healthcare	Retail	Manufacturing
Not evaluating	0%	2%	0%	0%	3%
Evaluating	43%	32%	21%	24%	33%
Piloting	34%	37%	32%	17%	19%
Limited production	18%	20%	19%	24%	24%
Scaled production	5%	9%	28%	35%	21%

Most organizations sit in the middle of the readiness scale. The readiness index averages **61–62%**, showing many teams are still closing gaps in the data foundation required for production. Uneven readiness shows up immediately in ROI expectations, as organizations that report stronger readiness are more confident that agentic AI spend will deliver durable returns.

Organizations that report being fully prepared are almost universally confident **(98%)**, saying they have **at least 70%** confidence that their data infrastructure will deliver the returns they expect. Among the least prepared, only **16%** report that level of confidence.

Confidence that data infrastructure can deliver ROI on AI investment by readiness segment

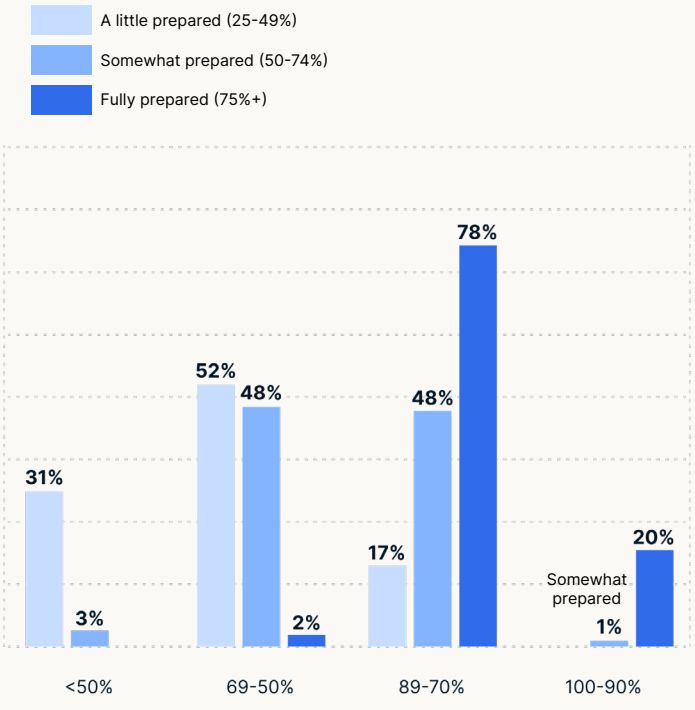

Readiness also shapes how broadly agentic AI is deployed. More prepared organizations are more likely to use agentic AI both for internal workflows and customer-facing products, while less prepared organizations more often keep it limited to internal operations and analytics (**57%**).

Primary approach to agentic AI by readiness segment

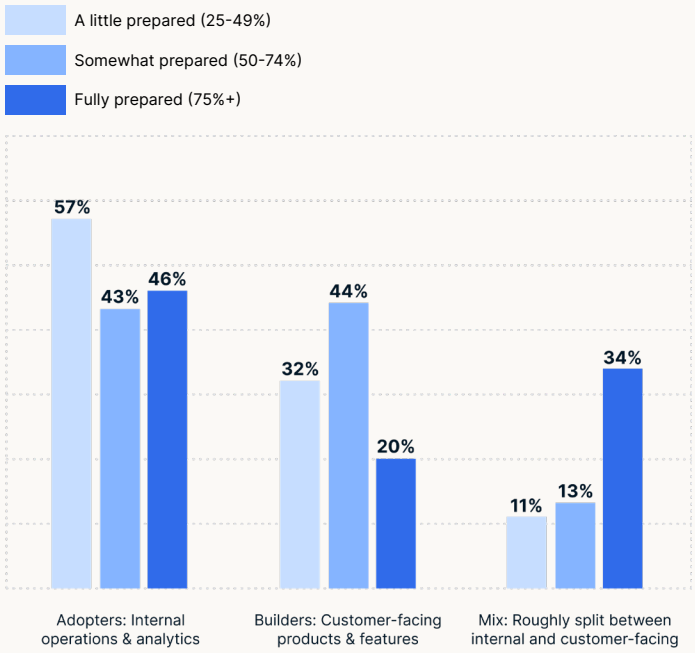

Keeping agentic AI limited to internal workflows can still create efficiency gains, but it also narrows the path to revenue-linked outcomes. Teams that deploy agents in customer-facing products have more direct ways to connect performance to metrics like retention, conversion, and growth.

The way teams define ROI suggests agentic AI is being evaluated less as a cost-cutting tool and more as a driver of customer and operational outcomes.

Organizations use 3 metrics on average to evaluate ROI:

- Customer experience and retention **(45%)**
- Financial cost savings **(43%)**
- Product development throughput **(43%)**

Headcount reduction **(25%)**, also a cited metric, is less commonly treated as a primary measure.

❖ KEY TAKEAWAY 2

Foundational data challenges are slowing agentic AI progress

As agentic AI moves from the experimentation phase into real workflows, the limiting factor becomes less about model capability and more about whether the data layer can support autonomy.

In practice, infrastructure sets the ceiling on outcomes, so model improvements cannot compensate for weak data movement, governance, or interoperability. Gartner concurs, predicting that, through 2026, organizations will abandon 60% of AI projects due to a lack of access to AI-ready data.

Data infrastructure is the primary blocker

Production adoption puts the data layer under stress. The most common obstacles to realizing agentic AI goals cited by participants fall into 3 categories:

- **42%** data quality and lineage issues
- **39%** sovereignty and regulatory compliance
- **39%** security and privacy risk

Other challenges, including strategy or resource gaps (**36%**) and skills shortages (**33%**), are cited less frequently as immediate blockers. Additionally, concerns like lock-in tend to emerge later, once organizations move from early adoption to scaling and long-term architecture decisions.

Strategy and staffing still matter, but many of those challenges become harder when teams cannot trust the data being fed into agentic systems. Agentic AI depends on clean, governed, well-documented data, and many organizations do not yet have that foundation in place.

Current large language models (LLMs) do not safely improvise around missing, stale, or non-compliant inputs. They can only make decisions and take actions based on the data they have access to. An agent powered by low-quality or poorly governed data does not get better over time; it simply makes incorrect decisions faster.

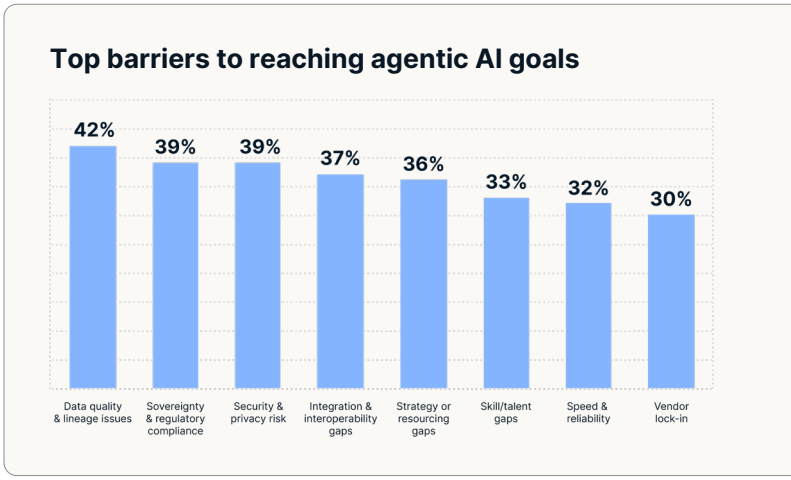

Governance and sovereignty now determine vendor selection

As agentic AI moves into production, often inside regulated workflows, data sovereignty and governance become decisive criteria for vendor access. Nearly two-thirds (**65%**) of organizations would either block or heavily restrict vendors who cannot meet governance and sovereignty standards.

- **25%** of organizations say they would decline a vendor outright if that vendor cannot meet data sovereignty and governance requirements.
- **40%** would only allow that vendor to handle non-sensitive or low-risk workloads.

In an agentic AI context, that matters because systems are being asked to act autonomously inside production environments and require strict control.

65% of organizations would either block or heavily restrict vendors who cannot meet governance and sovereignty standards.

❖ KEY TAKEAWAY 3

The interoperability imperative

Interoperability now shapes whether agentic AI can move beyond isolated pilots to profitability because agents only create value when they can work across the systems where data lives and decisions happen. Just as important, they need high-throughput access to that data.

For AI agents operating at scale, speed matters: relying solely on direct API calls for fast access is not a substitute for a strong, centralized data foundation. Without that foundation, organizations risk trading short-term speed for fragmentation, inconsistency, and limited scalability.

Platform-native agents often have more direct paths to data and actions inside their own ecosystems, while third-party agents are more often limited by APIs, sync schedules, and usage-based constraints.

Those differences affect not only what agents can access, but how quickly and reliably they can act on it. Over time, those differences create uneven agent performance and uneven risk, even when organizations use comparable models.

Importance of platform extensibility & interoperability to deliver on AI/data decisions

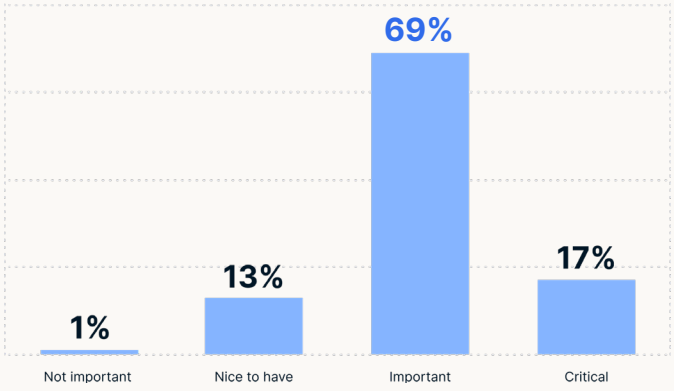

86% say platform extensibility and interoperability are important or critical to delivering AI and data-driven decisions.

In order to automate computer tasks, agentic AI depends on the ability to easily integrate with other tools, technologies, and platforms, as well as access to complete, clean, and reliable data. Poor interoperability hinders agentic AI's ability to operate systems or access necessary context.

Moreover, interoperability can create a 2-tier environment for agents. Platform-native agents often reach relevant data and actions more directly, while other agents are limited by interfaces and access constraints, resulting in uneven agent capability in production that's driven more by access than by model quality.

The systems agentic AI most commonly integrates with today are:

- Data warehouses and lakehouses (**42%**)
- Business applications like Salesforce and SAP (**40%**)
- BI and analytics tools (**38%)**

Agentic AI is most effective when it operates on centralized, governed, and continuously refreshed data. Data warehouses and lakehouses provide the foundation agentic systems need to reason, act, and adapt with confidence, bringing together operational, analytical, and historical data in a single environment.

For organizations looking to scale agentic workflows, data warehouses, and especially lakehouses, offer clear advantages:

- Improved data reliability
- Reduced fragmentation across source systems
- Faster access to up-to-date insights
- Open lakehouse formats that prevent lock-in

Getting data into the warehouse or lake requires additional consideration. When data pipelines are automated and resilient, teams can spend less time managing integrations and more time designing

intelligent agents that drive real business outcomes. As agentic AI matures, access to a trusted data foundation is emerging not as a nice-to-have, but as a core requirement.

Data integration is a perceived driver of vendor lock-in

Almost one quarter (**24%**) of organizations are extremely concerned that vendor lock-in will hinder their compliance and adaptability, with most others expressing at least moderate concern. For organizations that worry about lock-in, the data integration tools stand out as the biggest perceived risk (**34%**), ahead of cloud providers (**24%**) and major application vendors (**23%**).

As agentic AI becomes more tightly connected to how data is moved, governed, and refreshed, the integration layer becomes a strategic control point. When that control point is proprietary, it can effectively determine which agents get timely, complete access to trusted data and which agents are forced to operate with partial context. Choosing a closed or proprietary approach risks limiting future model adoption, constraining architectural flexibility, and increasing operational and compliance exposure.

Modern data integration strategies increasingly emphasize open data standards, such as Iceberg and Delta Lake conversion, to avoid vendor lock-in. By enabling interoperability across platforms, organizations can future-proof their data foundations for AI and evolving analytical needs.

CONCLUSION

What data leaders should do next: A blueprint for readiness

Organizations that are ready for agentic AI are building a resilient data foundation that can constrain, govern, and feed whatever AI comes next.

"The agentic AI readiness index" reveals consistent patterns among organizations that describe themselves as fully prepared:

- **They are more prepared across all 4 data foundation pillars**, with notable advantages in freshness and reliability (73% ready) and interoperability (64% ready).
- **Nearly all (98%) report at least 70% confidence** that their data infrastructure will deliver AI ROI, versus only 16% among the least prepared.

Their advantage is not model choice; it comes from infrastructure that keeps agentic AI governed and consistently fed with trusted data, which is what keeps risk and cost under control as adoption scales.

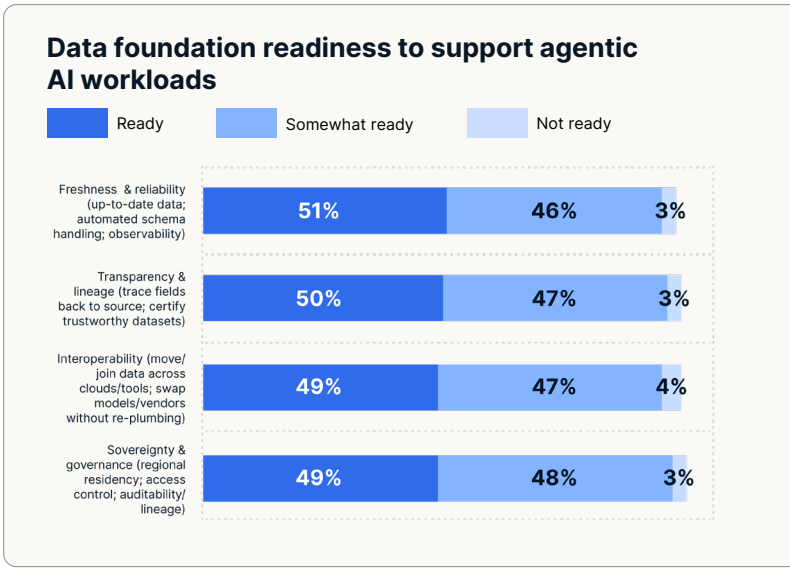

Closing the readiness gap comes down to 3 pillars — each one a deliberate investment in the data foundation agentic AI requires.

PILLAR 1:
Automate data movement to ensure freshness and reliability

Agentic AI decisions are only as good as the data available in the moment. Totally prepared organizations treat automated data movement as essential infrastructure, so agents always operate on fresh, reliable context as they move from pilots to production.

What prepared organizations do differently:

- Build **always-on** pipelines that refresh critical data continuously, not on batch schedules
- Use **automated schema handling** to reduce break/fix work as sources change

- Treat **observability** as standard instead of optional
- Standardize delivery into **central warehouses/ lakehouses** so agents have a single source of governed context

Less-prepared organizations absorb the cost of static data through lag, breakage, and ongoing break/fix work. Prepared organizations avoid that tax by keeping data movement automated, observable, and always-on so agents can operate on reliable context at scale.

PILLAR 2:
Enforce strong governance for trust and compliance

Agentic AI magnifies the consequences of weak governance. Because agents can operate at speed and scale, making decisions and triggering actions without human oversight at every step, gaps in governance can propagate quickly across systems and environments.

A viable agentic AI strategy requires enforceable governance mechanisms, including:

- **Access controls** that limit what agents can see and do
- **End-to-end lineage and auditability** to explain decisions, validate inputs, and support review
- **Regional sovereignty enforcement** that keeps data handling aligned with jurisdictional requirements
- **Designing for agent neutrality**, so any approved agent can access the same governed data through open interfaces

Organizations need enough confidence in these controls before allowing an AI system to interact with production environments with any degree of autonomy.

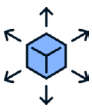

PILLAR 3:
Demand an open architecture for interoperability and anti-lock-in

Interoperability is emerging as a defining requirement for agentic AI success. As agents increasingly operate across data warehouses, business applications, and analytics tools, the ability to move, access, and combine data across systems determines whether agentic AI can scale beyond isolated pilots into core business workflows.

There is strong consensus among data leaders on this point:

- **34%** cite data integration as the single greatest source of lock-in risk — ahead of cloud providers (24%) and application vendors (23%).
- Totally prepared organizations are **significantly more likely** to consider interoperability 'critical' for AI decisions compared to the market average (29% vs. 17%).

This reflects a growing recognition that while infrastructure and applications may change, tightly coupled data pipelines are often the hardest — and most expensive — constraints to unwind.

A viable agentic AI strategy requires an Open Data Infrastructure that enables change without disruption, including:

- Choosing data integration platforms that support **open standards and broad connectivity**, rather than locking organizations into a single vendor or cloud ecosystem
- Building on a **central, governed, and interoperable data platform** to support agentic AI at scale
- Ensuring new models and AI services can be **added or replaced without rebuilding core data pipelines**

A vendor-neutral, automated data integration layer plays a critical role in this approach. By keeping data portable, governed, and continuously available across systems, it preserves flexibility, reduces lock-in risk, and provides a more durable foundation for agentic AI as technologies and requirements evolve.

Agentic AI pays off when the data foundation can support autonomy without adding operational or compliance risk. The organizations that describe themselves as fully prepared invest in the following capabilities: fresh and reliable data flows, transparency and lineage, enforceable governance, and open interoperability. Build those conditions, and you move toward the same endpoint they report: higher readiness, higher ROI confidence, and scaled deployment.

Methodology & demographics

This report is based on a survey of 400 data professionals across the United States, United Kingdom, EMEA, and Asia-Pacific. Respondents work in mid-sized and large enterprises and play an active role in data and analytics decision-making, including data architects, data engineers, data team leads, and analytics leaders responsible for building and operating agentic AI and the data infrastructure that supports it.

Participants come from data-intensive industries where automation and analytics are already central to day-to-day operations, with the largest groups drawn from technology, finance, healthcare, retail, and manufacturing.

To ensure the findings reflect established enterprise environments, the survey was limited to organizations with at least 2,000 employees in the US and EMEA and 500 or more employees in Japan, Australia, and Singapore. In total, the sample includes 100 respondents from the United States, 100 from the United Kingdom, 50 each from Germany, France, and Japan, and 25 each from Australia and Singapore, providing a view across major regulated and high-growth markets.