Use This Checklist to Ensure Your Data Is Ready for the Agentic AI Era

20 August 2025 - ID G00835373 - 11 min read

By: Deepak Seth, Roxane Edjlali

Initiatives:Data Management Solutions; Analytics and Artificial Intelligence; Architect, Implement and Scale Data and Analytics Solutions

Most organizations underestimate the additional data requirements for AI agents and agentic AI. This research guides data and analytics leaders to assess and enhance data readiness — ensuring right-time access, unified accessibility, contextual clarity, active metadata management and robust security to empower agentic AI.

More on This Topic

This is part of an in-depth collection of research. See the collection:

- How AI Agents Impact D&A Leadership, Governance and Data Management

Overview

Key Findings

- Many organizations underestimate the expanded data requirements and risks associated with deploying AI agents and agentic AI, leading to failed or underperforming initiatives.
- AI agents require data that meets clearly defined quality, timeliness, context and accessibility SLAs — established according to the specific demands of each AI use case — so they can operate effectively and autonomously at the necessary scale and speed.
- Insufficient metadata management, trust assessment, lineage tracking and unified access can result in agents making incorrect or untrustworthy decisions.
- Autonomous agents increase the risk surface — the total number of ways data can be misused or exposed — making robust governance, monitoring and security essential.

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

Recommendations

- Implement rigorous, use-case-specific data readiness checks for every AI agent or agentic workflow to prevent errors — such as deploying agents outside their intended context or overlooking data drift — that can undermine reliability and outcomes.
- Align data delivery infrastructure to the latency, frequency and autonomy needs of agentic AI by investing in integrated platforms that provide seamless, secure access to all relevant data sources.
- Enhance metadata management and data lineage practices to ensure agents have the context and provenance needed for accurate, trustworthy decisions.
- Strengthen data governance, access controls, observability and automated governance to secure contextual data governance and mitigate new risks introduced by autonomous systems.

Strategic Planning Assumptions

By 2028, 33% of enterprise software applications will include agentic AI, up from less than 1% in 2024.

By 2028, at least 15% of day-to-day work decisions will be made autonomously through agentic AI, up from 0% in 2024.

By 2028, 40% of agentic AI projects will be canceled due to escalating costs, unclear business value or inadequate risk controls.

Introduction

As organizations accelerate their adoption of AI agents and agentic AI, many are discovering that traditional data management practices are no longer sufficient. The promise of agentic AI — systems capable of making decisions and executing complex workflows autonomously — depends fundamentally on the quality, accessibility and contextual richness of the underlying data.

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

There has been a 265% increase in venture capital (VC) investments in this space in 1Q25 from 4Q24, according to PitchBook. Agentic AI conversations have exploded at Gartner. Gartner saw a 750% increase in AI-agent-related inquiries between the second and fourth quarters of 2024. [1] However, most of these are experiments or POCs at this stage and are mostly driven by hype. Consequently, organizations run the risk of greatly underestimating the cost and complexity of deploying AI agents and end up severely underdelivering on overhyped expectations.

According to Gartner client inquiries, many organizations piloting agentic AI report or fear unexpected challenges related to data readiness, including issues with data latency, fragmentation, lack of context and insufficient governance. These challenges not only impede the realization of business value from AI investments but also introduce new operational and regulatory risks — such as data breaches, unauthorized access or misuse of sensitive information — which can lead to violations of existing data protection and privacy regulations like GDPR, CCPA or sector-specific compliance requirements. As agents gain autonomy, organizations must ensure their AI practices align with established regulations to avoid legal and reputational consequences.

The approach outlined in this research is designed for data and analytics leaders, CIOs and business stakeholders, whose decisions and oversight are critical to enabling or scaling agentic AI across the organization. This checklist is most effective in environments where AI agents are expected to operate with increasing independence by making decisions, interacting with diverse data sources and executing tasks with minimal human oversight.

However, this approach may be less relevant for organizations implementing narrowly scoped, rule-based automation, regardless of their overall AI maturity. In such cases, traditional data management practices may suffice. For organizations seeking to harness the full potential of agentic AI, assessing and strengthening data capabilities is not optional — it is foundational to success.

Analysis

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

Unlock the Full Potential of AI Agents and Agentic AI

AI agents and agentic AI mark a significant evolution from traditional machine learning, shifting from systems that require human interpretation to those capable of autonomous, context-aware decision making and workflow orchestration. As organizations move from conventional AI to agentic solutions, data readiness requirements become more dynamic and complex — demanding not just periodic, batch-updated data but continuous, right-time and semantically consistent access across multiple systems.

Organizations may deploy solutions ranging from simple, rule-based agents to sophisticated agentic AI (see Quick Answer: What Are the Emerging Use Cases for AI Agents in the Financial Services Industry?). Data requirements and risks grow as autonomy increases. Traditional AI and machine learning models often operate on data that is periodically updated in batches — both during model development and in production — relying on static snapshots rather than real-time information.

In contrast, AI agents require input data that is updated more frequently, with integrated access across multiple systems and improved interoperability between data sources, enabling them to make timely and relevant decisions. Agentic AI raises the bar even further, demanding continuous, right-time data delivery and access to information that is semantically consistent with the agent's needs — including robust metadata, governance and adherence to defined SLAs. This enables agents to interpret business meaning, track data lineage and apply the appropriate context to each decision, supporting more reliable, transparent and autonomous operations. While agents were envisaged as orchestrators — consuming and acting upon data — emerging architectures blur the line, with advanced agents now capable of initiating actions and even producing new data or outputs.

To unlock the full potential of AI agents and agentic AI, organizations must further build upon their approach to AI-ready data using Gartner's established AI-ready data framework (see Quick Answer: What Makes Data AI-Ready?) and the 30 foundational preparation tasks listed in How to Evaluate AI Data Readiness. The data and tasks should then be contextualized for the unique demands of AI agents and agentic AI, addressing new dimensions — such as right-time data delivery, dynamic context, autonomous access, continuous quality monitoring, computational governance and explicit agent context — tailored to the specific requirements of each AI use case, rather than relying on one-size-fits-all standards from previous AI implementations. The following checklist will help you achieve these objectives.

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

Data Readiness Checklist for AI Agents and Agentic AI

The five steps outlined in the checklist provide a structured approach for ensuring data readiness, reliability and governance as agentic capabilities advance. For detailed definitions, distinctions, and assessment frameworks for AI agents and agentic AI, see Tool: AI Agent Assessment Framework. For a companion checklist that distills the guidance from the research to help you assess the data readiness for implementing AI agents or agentic AI, select the download below.

Download Data Readiness Checklist for AI Agents and Agentic AI

Guidelines for using the checklist to ensure data readiness for AI agents are as follows:

- **Understand the purpose:** *Use the checklist to evaluate and enhance your data's readiness for both AI agents and agentic AI, focusing on strategic alignment, operational efficiency, data quality, infrastructure, timeliness and data variety.*
- **Conduct a comprehensive review:** *Assess your current data systems and practices with input from IT, data management and business stakeholders for a holistic evaluation.*
- **Identify gaps and opportunities:** *Compare your current capabilities and prioritize improvements based on your organization's AI maturity and goals.*
- **Develop an action plan:** *Create a detailed plan to address identified gaps. Set clear objectives, timelines and responsibilities for achieving readiness.*
- **Monitor and update:** *Revisit the checklist regularly to track progress and adjust as technology and business needs evolve.*
- **Leverage insights for business decisions:** *Use the results to inform your AI strategy and ensure your organization is well-positioned to capitalize on both AI agents and agentic AI advancements.*

Step 1: Conduct Data Readiness Checks for Each Agent

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

Every AI agent should undergo a data readiness evaluation. This check ensures the agent can reliably access, interpret and utilize the specific data required for its intended functions and autonomy level. Organizations must determine when periodic or assumption-based validation suffices and when continuous, automated (such as checksum-style) validation is required. This prevents agents from acting on stale, incomplete or corrupted data, which could undermine business outcomes and erode trust in autonomous systems (see How to Evaluate AI Data Readiness).

Step 2: Align Data Delivery Capabilities to the Required Level of Autonomy

To support both basic AI agents and advanced agentic AI solutions, data management infrastructure must deliver information at the "right time" for each use case. "Right-time" data means delivering data with the latency, frequency and freshness required by the agent's operational context — ranging from periodic batch updates to real-time, event-driven streams for highly autonomous systems.

Establishing these requirements starts with a cross-functional analysis of the agent's tasks, decision life cycles, risk tolerance and business impact. Business stakeholders, data owners and technical teams should collaborate to define precise SLAs for data freshness, consistency, accessibility, security and cost-effectiveness.

For detailed implementation guidance, see 2025 Strategic Roadmap for the Data Fabric Architecture.

Step 3: Build Unified Data Accessibility That Scales With Agent Sophistication

Seamless access to diverse data sources — structured and unstructured, internal and external — is essential for AI agents at all levels of sophistication. Organizations should invest in unified data platforms and APIs, such as data fabrics, data virtualization layers or cloud-based data marketplaces, to provide scalable, secure and governed access to data (see Data Lakehouses Are a Key Strategic Pillar of Data Fabrics and Ecosystems). These platforms should support fine-grained access controls, robust authentication, audit trails, and monitoring to ensure both security and compliance.

Implementing a semantic layer — either within the data catalog or as an additional abstraction — helps provide consistent business definitions and relationships (see Rethink Semantic Layers to Support the Future of Analytics and AI).

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

However, consistent semantics alone is not sufficient. To fully meet the data requirements of agentic use cases, organizations must also ensure that data provenance, trustworthiness, quality, governance and security are maintained and made transparent to agents. This may involve leveraging knowledge graphs, business rules and advanced metadata management to provide agents with the context needed to accurately select, interpret and apply data in line with business intent and regulatory requirements (see How to Build Knowledge Graphs That Enable AI-Driven Enterprise Applications and State of Metadata Management: Aggressively Pursue Metadata to Enable AI and Generative AI).

Step 4: Ensure Contextual Clarity and Data Provenance

To perform effectively, AI agents must understand the context, lineage and meaning of the data they consume. Robust metadata management, data catalogs and lineage tracking are essential to ensure agents can interpret data definitions, relationships and trustworthiness (see Quick Answer: What Is Active Metadata?). As agentic solutions advance, organizations should enhance their data architecture with knowledge graphs and contextual metadata (such as data domain, the presence of client data, IP data, etc.) for unstructured data that links data assets to business processes and goals (see Develop Unstructured Data Management Capabilities to Support GenAI-Ready Data). This enables agents to reason about cause and effect, personalize actions, and make autonomous decisions when safe to do so. Each agent should also express its operational context explicitly and publish its relationships to the knowledge graph if it is being used, ensuring transparency and discoverability of its capabilities and dependencies.

Step 5: Manage and Secure Data for Agentic Workflows Actively

As AI agents and agentic AI solutions become more autonomous, the risk of data misuse, leakage and manipulation increases. Organizations must shift from static, periodic controls to continuous, automated and code-driven governance — known as automated governance (see Hype Cycle for Data and Analytics Governance, 2025). This involves implementing multilevel guardrails — such as granular access controls, policy-based usage restrictions, automated data quality checks, real-time monitoring and audit logging — as integral parts of agentic workflows. Ongoing monitoring and automated validation ensure that all data — including sensitive but nonpersonal information like intellectual property and confidential strategies — is used securely, ethically, and compliantly with regulatory and corporate requirements. Continuous, checksum-style governance is essential to maintain data reliability and compliance as agents act on behalf of the enterprise (see Six Top Practices to Assure That Your Unstructured Data Is AI Ready).

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

Risks of Inadequate Data Readiness

If context and knowledge graphs are lacking, agents may misinterpret data, producing erroneous or untraceable results. Without consistent, ongoing data readiness assessment for every agent invocation, errors and poor decisions are likely to proliferate. The cumulative risk grows when multiple agents operate without provable data readiness or are used outside their intended context. Furthermore, the lack of standardized protocols for managing context — both client- and server-side — exposes organizations to even greater risks as agentic adoption accelerates.

Contributors

Mark Beyer

Evidence

This research is based on client inquiries and prior Gartner research.

[1] Emerging Tech: Avoid Agentic AI Failure: Build Success Using Right Use Cases

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

When to Use or Not to Use AI Agents

Ready Your Data for AI

6 Case Studies To Help Navigate Ethical AI Dilemmas

A Journey Guide to Manage AI Governance, Trust, Risk and Security

This research note is restricted to the personal use of SAM@BLUECONIC.COM.

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by Gartner's Usage Policy. Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "Guiding Principles on Independence and Objectivity." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

This research note is restricted to the personal use of SAM@BLUECONIC.COM.