Emerging Tech: Design AI Agents to Work With Messy Data

12 September 2025 - ID G00827266 - 10 min read

By: Alfredo Ramirez IV, Aakanksha Bansal, Anushree Verma, Danielle Casey, Akhil Singh

Initiatives:Emerging Technologies and Trends Impact on Products and Services

> Successful use of agentic AI patterns does not require perfect data for all use cases. Product leaders must build agentic AI systems architected to accomplish work using imperfect, incomplete, and disparate data, similar to how humans work, or else risk losing business and adoption to solutions that do.

Overview

Key Insights

- Impactful, high-value agentic AI systems in production demonstrate that C-level leaders must not wait on large-scale data initiatives when they can gain organizational benefits from agentic AI now by carefully selecting imperfect data-tolerant use cases.
- When agentic AI solutions are designed to process data in situ to meet fit-for-purpose use cases, the barriers to adoption are lowered and time-to-value can be accelerated.

Recommendations

Product leaders seeking to accelerate agentic AI solution adoption must:

- Make ROI discussions tangible by going to market demonstrating how your agentic AI solution augments human work using data within existing operational workflows.
- Focus product development, agent design, and context engineering on leveraging in situ data that is representative of the use case and use comprehensive test plans that assume operational data will be incomplete, disparate, or missing across systems.

Analysis

This document was revised on 18 September 2025. The document you are viewing is the corrected version. For more information, see the Corrections page on Gartner.com.

About This Document

This document is part of Gartner's case-based research (CBR) into the current state and future direction of Agentic AI. This project involved outreach to 36 tech providers, completing interviews across *27* agentic AI providers globally and analyzing over *170* adopter use-case studies over a period of four months starting in January 2025.

Trend Description

Using data "in situ" refers to using data where it already resides and in its current state. This is how humans accomplish much of the digital work they do today. Data that a worker needs to triage an alert, answer a customer query or resolve a ticket always exists in multiple disparate systems and requires domain-specific knowledge, gathering context and pattern recognition to come to a conclusion.

This is important because agentic systems can accomplish work in the same way if they are given domain-specific knowledge, access to tools to fill in missing information, and prompted to achieve the proper context before coming to a conclusion.

The following is a nonexhaustive list of vendors bringing agentic solutions to market that are able to accomplish real-world digital labor, like cybersecurity threat hunting, expense auditing, first-line technical support tasks, etc. using data in situ: Airrived; Crogl; Exaforce; Glean; Lyzr; Relevance AI; ReliaQuest; Tuskira

Gartner's position is that AI-ready data means that your data must be representative of the use case, of every pattern, error, outlier and unexpected emergence that is needed to train or run the AI model for the specific use. Data readiness for AI is not something that can be established once and for all, nor can you build ahead of time for all your data. It is a process and a practice based on the availability of metadata to align, qualify and govern the data

Agentic AI is an approach to building AI solutions based on the use of one or multiple software entities classified, completely or at least partially, as AI agents. AI agents are autonomous or semiautonomous software entities that use AI techniques to perceive, make decisions, take actions, and achieve goals in their digital or physical environments.

Operationally useful agentic AI systems are grounded in domain specific knowledge and organizational context that aids them in finding data that they need to build their own context and come to conclusions. Vendors observed in Gartner CBR research and inquiry, that are successfully deploying systems for messy data-tolerant user cases using organizational data as it is and where it is, expertly utilize system design that respects the "fundamental trade-off between certainty and scope in symbolic and generative AI." [1] This requires domain-specific knowledge and strategic use of deterministic and probabilistic AI components, as shown in Figure 2.

Product leaders must build agentic solutions grounded in domain-specific knowledge and QA tested with the assumption that fit-for-purpose data available for task execution will be disparate, incomplete, or missing across the environment in which the agents operate.

Figure 1. Design AI Agents to Work with Messy Data

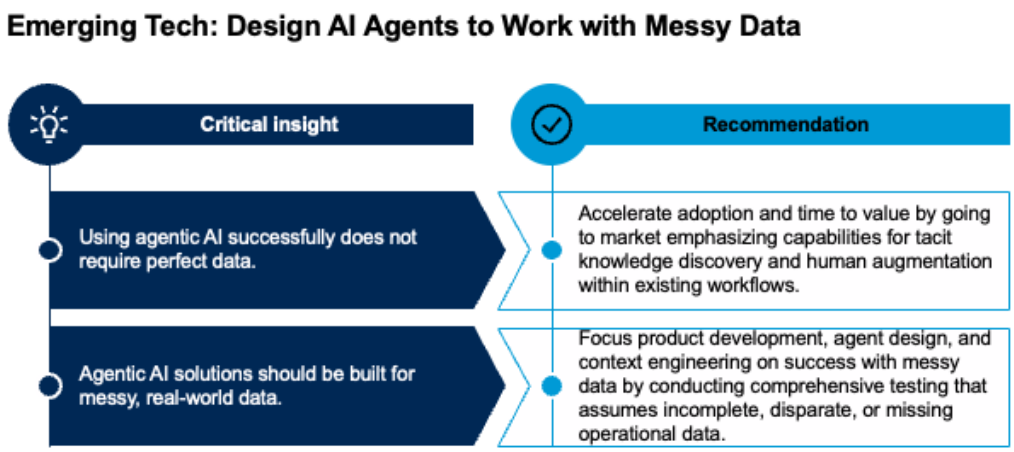

Source: Gartner
827266

Gartner

Gartner 2025

Critical Insight: Using agentic AI successfully does not always require perfect data.

Near-Term Implications for Product Leaders

- Most organizations struggle with incomplete, disparate and unorganized data, and agentic systems designed to deal with this reality lower the adoption barrier, as well as the time to realize agent value.
- Organizational leaders do not always need to advocate for upfront data infrastructure redesign to cultivate broad sets of curated data as a prerequisite for targeted agentic AI deployment.
- There is a significant opportunity to accelerate time-to-value and demonstrate tangible ROI much faster by focusing on designing domain-specific agentic AI solutions that can use existing operational workflows.

Recommended Actions for the Next Six to 18 Months

- Accelerate agentic AI adoption and time to value by going to market emphasizing your agentic AI solutions' capabilities for tacit knowledge discovery, that is capturing and leveraging organizational expertise by identifying patterns, context, and implicit relationships from available structured and unstructured data, and how it lowers a major barrier to entry for your prospects' agentic AI adoption plans.
- Help your prospect champions build a strong business case for your solutions by working with them to quantify organizational savings through agentic system implementation.
- Demonstrate tangible ROI by showcasing how your agentic AI solution augments human work within existing operational workflows.

Critical Insight Analysis

Agentic AI solutions are capable of planning, inferring, and taking action based on existing data and tool usage. Crucially, these new design patterns do not require transforming legacy data infrastructures. Instead, they integrate into the current architecture to extract value from data where it already persists.

This means agentic solutions tap directly into operational databases, ticket histories, unstructured data, policy and organizational documentation, as well as interviewing human subject matter experts (SMEs) to reconstruct tacit organizational knowledge, operational context and carry out tasks. Successful vendors also demonstrated iterative agentic output validation via other agents. An agent conducting validation is powered by a different goal and context and may as well be using a different model to conduct its analysis. See Figure 2 item 9 as well as Table 1 item 9 "probabilistic guardrails."

However, these benefits are not being communicated well to the market and this represents an opportunity for product leaders in this space to enhance messaging to prospects. Many of whom are weary of yet another data-heavy project before being able to take advantage of AI technology.

Critical Insight: Agentic AI solutions should be built for messy, real-world data.

Near-Term Implications for Product Leaders

- Our research found that agentic output evaluation is key. Vendors like Lyzr and Relevance AI enable multiagent systems where outputs of one agent can be checked by another whose goal is to test for things like coherence and alignment with the original request. (See: Architecture to Build AI Agents - multiagent modularity pattern)

- As agentic AI frameworks mature, there will be a growing demand for more robust multi-agent systems that communicate seamlessly, share persistent context, and achieve complex business objectives collaboratively with humans.

Recommended Actions for the Next Six to 18 Months

- Focus product development, agent design, and context engineering on scenarios that involve messy data. Conduct comprehensive QA testing in situations with incomplete, disparate, or missing data that mimics organizational realities.

- Utilize agentic output evaluation where outputs of one agent can be checked by another whose goal is to iteratively test for things like coherence, alignment and relevance for the original request.

- Accelerate R&D efficiencies in this area by leveraging other disruptor technologies like hypersynthetic data and intelligent simulation — e.g., a simulation twin that mirrors the intended solution's use cases and using multiple LLMs to simulate employees' behaviors in that context. (See: Emerging Tech Impact Radar: Intelligent Simulation)

Critical Insight Analysis

A GenAI system for processing messy data must have a semantic model that is able to contextualize even incomplete information it has access to. Agentic AI brings together GenAI capabilities that can make sense of messy and inconsistent data using:

- Retrieval-augmented generation (RAG) for contextual grounding
- Tool access
- Human input
- Domain specific language models

These capabilities are combined with both deterministic and probabilistic guardrails, such as reliability/reviewer agents (also known as guardian agents) for output evaluation, to create systems that are adaptive and reliable.

> **Agentic systems are most useful when implemented using composite AI, multiple focused subsystems, grounded in organizational knowledge, with access to humans in the loop, augmented by domain-specific models, and with outputs controlled by deterministic and probabilistic guardrails.**

Taken together, these systems represent a way for product leaders to provide immense value to end users if the product is designed to take advantage of agentic benefits, and those benefits are communicated early and often to ease apprehensions around adopting AI agents.

Figure 2 is a simplified example of an agent defined by a goal and powered by a set of subprocesses to accomplish digital work using existing data in existing systems. The accompanying Table 1 describes each example component.

Successful vendors created agentic systems that were resilient in similar architectures, allowed for human input as shown in workflow steps No. 4 and No. 4a, and where outputs were iteratively validated via deterministic (static policies) and probabilistic (other agents) guardrails as in steps No. 8 and No. 9.

Figure 2. Example Agentic Design

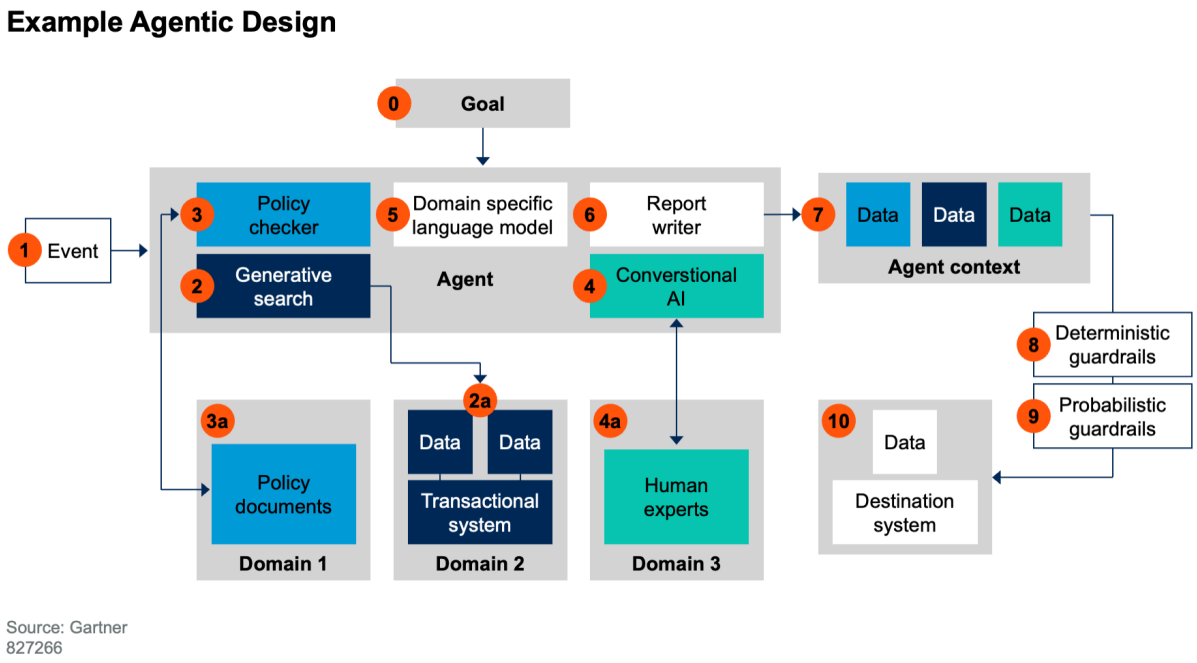

Source: Gartner
827266

Gartner

Gartner 2025

Table 1: Legend for Example Agentic Design

(Enlarged table in Appendix)

#	Agentic Component	Description
0	Goal	Part of the agent configuration and used to direct planning and outputs.
1	Event	Ecosystem event or other input, human or otherwise, signals agent work to start.
2	Generative Search	Based on the Goal and Event, this can be used to gather context from tools and historical data that the agent has access to.
2a	Transactional System	One or more downstream tools and systems to be queried or utilized to establish context en route to the Goal.
3	Policy Checker	Grounded on unstructured data to ensure outputs meet organizational policies.
3a	Policy Documents	Unstructured grounding data using RAG to add organizational context, akin to how a human might consume it.
4	Conversational AI	NLP component for bidirectional communication between the agent and humans that have useful context to achieve the Goal.
4a	Human Experts	To be contacted over standard channels like instant messaging and email to provide context to the agent, similar to a coworker interaction.
5	Domain Specific Language Model (DSLM)	A smaller-parameter model is often trained on a particular area, like law, accounting, or cybersecurity, to ensure more accurate and reliable outputs.
6	Report Writer	Summarizes outputs based on context, which may include Policy Documents outlining preferred organizational formatting.
7	Agent Context	This represents the readied output based on the organizational context gathered by the agent from disparate sources No. 2-4 and augmented by domain knowledge from No. 5.
8	Deterministic Guardrails	Rules to avoid unwanted outcomes and improve reliability include static policies, bias and toxicity checks, redaction, acceptable use, etc. If this check fails, the response will be rerouted to the agent to try again.
9	Probabilistic Guardrails	This represents an agent or agents with a goal of safety, quality assurance, or compliance, acting as a dynamic check in addition to deterministic guardrails. If this check fails, the response would be rerouted to the agent to try again (see Guardians of the Future: How CIOs Can Leverage Guardian Agents for Trustworthy and Secure AI).
10	Destination System	Downstream system(s) that receive the now aligned, enriched, contextualized, and quality-assured agentic response.

Source: Gartner (September 2025)

Acronym Key and Glossary Terms

AI	artificial intelligence
agentic AI	A type of AI system characterized by autonomous decision-making, multistep execution, and real-time adaptation, capable of planning, reasoning, and taking action based on present data
in situ	In its original place or position, it refers to leveraging data where it already resides without a complete overhaul of existing data architecture
LLM	large language model
RAG	retrieval aAugmented generation
ROI	return on investment

Evidence

[1] A Conjecture on a Fundamental Trade-Off Between Certainty and Scope in Symbolic and Generative AI, Springer Nature Link.

Recommended by the Authors

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

Emerging Tech: Hyper-Synthetic Data Unleashes AI Innovation by Avoiding Shortfalls of Real Data

Emerging Patterns for Building LLM-Based AI Agents

© 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.

Table 1: Legend for Example Agentic Design

#	Agentic Component	Description
0	Goal	Part of the agent configuration and used to direct planning and outputs.
1	Event	Ecosystem event or other input, human or otherwise, signals agent work to start.
2	Generative Search	Based on the Goal and Event, this can be used to gather context from tools and historical data that the agent has access to.
2a	Transactional System	One or more downstream tools and systems to be queried or utilized to establish context en route to the Goal.
3	Policy Checker	Grounded on unstructured data to ensure outputs meet organizational policies.
3a	Policy Documents	Unstructured grounding data using RAG to add organizational context, akin to how a human might consume it.
4	Conversational AI	NLP component for bidirectional communication between the agent and humans that have useful context to achieve the Goal.
4a	Human Experts	To be contacted over standard channels like instant messaging and email to provide context to

		the agent, similar to a coworker interaction.
5	Domain Specific Language Model (DSLM)	A smaller-parameter model is often trained on a particular area, like law, accounting, or cybersecurity, to ensure more accurate and reliable outputs.
6	Report Writer	Summarizes outputs based on context, which may include Policy Documents outlining preferred organizational formatting.
7	Agent Context	This represents the readied output based on the organizational context gathered by the agent from disparate sources No. 2-4 and augmented by domain knowledge from No. 5.
8	Deterministic Guardrails	Rules to avoid unwanted outcomes and improve reliability include static policies, bias and toxicity checks, redaction, acceptable use, etc. If this check fails, the response will be rerouted to the agent to try again.
9	Probabilistic Guardrails	This represents an agent or agents with a goal of safety, quality assurance, or compliance, acting as a dynamic check in addition to deterministic guardrails. If this check fails, the response would be rerouted to the agent to try again (see Guardians of the Future: How CIOs Can Leverage Guardian Agents for Trustworthy and Secure AI).
10	Destination System	Downstream system(s) that receive the now

	aligned, enriched, contextualized, and quality-assured agentic response.

Source: Gartner (September 2025)