

Five Ways Healthcare AI Agents Fail Without Clean Patient Identity Data

AI in Healthcare: Why Trusted Identity Is the Foundation for Accuracy, Explainability, and Scalability.



Executive Summary

Healthcare organizations are rapidly exploring AI agents to improve productivity, reduce administrative burden, and automate increasingly complex clinical workflows. As these systems become more autonomous, the quality of the identity data they rely on becomes critical to their success.

AI agents are only as effective as the data they can access. For healthcare AI that engages patients, providers, consumers, or organizations, trusted identity is the starting point. When those records are fragmented across disconnected systems, AI agents inherit incomplete and inconsistent data. The result is inaccurate recommendations, failed automations, governance challenges, and reduced trust in AI outcomes.

Enterprise master data management (MDM) provides the trusted identity foundation AI agents require. By resolving fragmented records, enriching identity data, verifying identities, and providing transparency into identity decisions, MDM enables AI that is accurate, explainable, and scalable.

This whitepaper examines five common ways AI agents fail without clean identity data and explores how modern identity intelligence helps healthcare organizations build a stronger foundation for enterprise-scale AI.



Key Takeaways

- **AI initiatives require AI-ready data foundations.**
Poor data quality remains one of the primary barriers to enterprise AI adoption; **60% of AI projects will be abandoned through 2026 due to a lack of AI-ready data.**¹
- **Poor identity data creates measurable business risk.**
Poor data quality costs organizations an average of **\$12.9 million annually**, while also degrading AI performance and increasing operational inefficiencies.²
- **Context informs meaningful AI outcomes.**
Social determinants of health and other non-clinical factors significantly influence health outcomes, making identity enrichment critical for AI systems that require a more complete understanding of the individual.³
- **Transparency is essential for enterprise AI governance.**
Organizations need visibility into identity resolution and verification decisions to confidently govern AI-driven actions and meet increasing expectations for accountability.
- **Trusted identity powers agentic AI.**
Enterprise MDM resolves fragmented identities and provides the identity resolution, enrichment, verification, and explainability needed for AI agents to operate accurately, consistently, and at enterprise scale.



of AI projects will be abandoned through 2026 due to a lack of AI-ready data



Poor data quality costs organizations the equivalent of 12.9 million dollars annually.

AI Agents Require Trusted Identity Data

Unlike traditional analytics, AI agents do not simply generate insights. They execute workflows and take action on behalf of healthcare organizations. When identity data is fragmented or inaccurate, AI agents amplify those weaknesses, increasing the risk of incorrect recommendations, operational errors, and reduced trust. A misidentified patient is no longer just a data quality issue. It becomes an AI reliability issue. Many organizations are investing aggressively in agentic AI, yet only **21% of enterprises report having a mature governance model for it.**⁴ As autonomous systems become more prevalent, establishing trusted identity data is becoming a prerequisite for successful AI adoption.

Only 21% of enterprises report having a mature governance model for AI.



The Missing Foundation Behind Reliable AI

AI agents depend on trusted identity data to execute workflows across complex healthcare environments. Yet patient, provider, and organizational identities are often fragmented across disconnected systems, resulting in duplicate records, incomplete information, and inconsistent views of the same individual or entity. Enterprise MDM provides the trusted identity layer that unifies records across systems, creates a complete view of individuals and organizations, and enables AI agents to operate with enterprise-grade identity intelligence.



Five Ways AI Agents Fail Without Clean Identity Data

FAILURE #1

Identity Mismatching Gives AI Agents an Incomplete Patient View



The Challenge

The most consequential AI agents begin with the same fundamental question: Who am I acting on? When patient identities are fragmented across multiple systems, AI agents absorb duplicate records, conflicting information, and incomplete histories. As a result, recommendations, workflows, and automated actions may be based on an inaccurate understanding of the individual.



Why It Happens

Healthcare organizations manage identity data across numerous clinical, operational, consumer, and partner systems. Without an enterprise identity foundation, these systems create fragmented records that AI agents cannot reliably reconcile on their own.



The Opportunity: Establish a Trusted Golden Record

Enterprise MDM creates a persistent golden record by connecting identity data and accurately resolving identities at scale. This provides AI agents with a trusted, unified view of each individual before any action is taken.

As Gartner notes, **60% of AI projects will be abandoned through 2026 due to a lack of AI-ready data, highlighting the importance of establishing an enterprise identity infrastructure before scaling AI initiatives.**¹



Outcome

With a trusted golden record, AI agents operate from a complete and accurate understanding of the individual, enabling more reliable automation and more informed decision-making.

FAILURE #2

Contextual Blindness Leaves AI Without the Full Patient Picture



The Challenge

AI agents can only make informed decisions when they have a complete view of the individual. Yet many healthcare systems lack critical demographic, behavioral, social, and environmental information that provides the context behind clinical data. Without that broader perspective, AI recommendations may be incomplete or less relevant.



Why It Happens

Contextual data is often scattered across disconnected systems or unavailable altogether, with each application maintaining only a partial view of the individual. AI agents' ability to understand the full picture becomes limited.



The Opportunity: Add Intelligence Through Identity Data Enrichment

Enterprise MDM creates a unified identity foundation that can be enriched with trusted external data, providing AI agents with more complete context for decision-making. An MDM solution with native identity enrichment capabilities helps healthcare organizations build a more comprehensive understanding of the individuals they serve.

The need for this broader context is significant. Social determinants of health influence up to **80% of health outcomes, yet collecting accurate and complete, individual-level SDOH data remains a persistent challenge.**^{3,5} As AI agents increasingly support clinical and operational decisions, identity enrichment helps bridge these information gaps.



Outcome

With richer identity data, AI agents can deliver more informed recommendations and better support decisions across healthcare workflows.

FAILURE #3

Training Drift Erodes the Accuracy of AI Insights



The Challenge

AI models are only as reliable as the data used to train and inform them. Duplicate identities, outdated information, incomplete records, and inconsistent data introduce errors that reduce the accuracy of AI-driven insights, automation, and decision-making.



Why It Happens

Traditional identity matching approaches struggle to keep pace with the complexity of modern healthcare ecosystems. As data moves across multiple systems, identity inconsistencies accumulate faster than manual processes can resolve them.



The Opportunity: Resolve Identity Before Data Reaches AI

Enterprise MDM improves data quality at the source by resolving fragmented identities before they reach AI models. Leading healthcare organizations require an enterprise MDM solution capable of accurately linking records despite incomplete or inconsistent data, creating cleaner, more reliable datasets for analytics and AI applications.

The business impact of poor data quality is substantial. Gartner estimates that organizations lose an average of **\$12.9 million annually due to poor data quality, making trusted identity data a requirement for enterprise AI.**²



Outcome

By improving identity data before it reaches AI systems, organizations enable more accurate models, more reliable automation, and better-informed decisions.



FAILURE #4

Verification Gaps Expose AI Interactions to Privacy, Security, and Fraud Risks



The Challenge

As AI agents increasingly interact with patients, providers, and consumers, they must be able to verify they are engaging with the correct individual. Without trusted identity verification, automated actions can introduce privacy, security, and fraud risks.



Why It Happens

Many organizations implement identity verification separately across individual applications and channels. Without connecting verification to a trusted enterprise identity, AI agents lack a consistent way to validate who they are interacting with.



The Opportunity: Anchor Verification to Identity Resolution

Enterprise MDM establishes an identity foundation that can be extended with real-time identity verification. An effective MDM solution integrates identity verification directly with enterprise identity data, enabling AI agents to interact with the correct individual across digital and clinical workflows while improving the reliability of enterprise automation.

The need for stronger verification continues to grow. Javelin Strategy & Research estimates that identity fraud generates more than **\$40 billion in losses**, underscoring the importance of trusted identity as organizations expand AI-driven interactions.⁶



Outcome

By combining trusted identity with identity verification, healthcare organizations enable AI agents to automate interactions more securely while reducing privacy, security, and fraud risks.



FAILURE #5

The Black Box Dilemma Makes AI Decisions Harder to Validate



The Challenge

As AI agents take on more responsibility across healthcare workflows, organizations must be able to understand and govern the decisions they make. Without visibility into how identities are resolved, it becomes difficult to validate AI-driven actions, support audits, or establish trust in automated outcomes.



Why It Happens


Identity matching often occurs behind the scenes with little transparency into why records were linked or how identity decisions were made. This lack of explainability creates governance challenges as AI adoption scales.



The Opportunity: Build Explainability into the Identity Layer

Enterprise MDM can provide visibility into identity-resolution decisions through match rules, scores, and stewardship workflows. More advanced explainability capabilities help data stewards and auditors to understand why records were linked, supporting stronger governance and greater trust in AI-driven workflows.

The National Institute of Standards and Technology (NIST) identifies explainability, interpretability, accountability, and transparency as core characteristics of trustworthy AI, underscoring the importance of making identity decisions understandable rather than opaque.⁷




Outcome

By embedding explainability into the identity foundation, healthcare organizations can govern AI agents more effectively while supporting greater accountability, transparency, and trust.

MDM Provides the Identity Capabilities for Agentic AI

The five AI agent failure points explored in this paper are different manifestations of the same underlying challenge: fragmented enterprise identity. Modern MDM brings together identity resolution, enrichment, data stewardship, governance, explainability, and activation into a single enterprise capability. Rather than solving individual identity problems one application at a time, it creates a shared identity layer that every AI agent, workflow, and system can rely upon.

As organizations expand agentic AI across clinical, operational, and consumer workflows, identity can no longer be treated as an application-specific function. It must become an enterprise capability. MDM provides that capability, enabling AI agents to operate consistently across systems, workflows, and use cases.



Verato MDM Cloud™ is the identity foundation

Verato MDM Cloud™ delivers the enterprise identity capabilities required to operationalize agentic AI at scale. **Referential Matching®** resolves identities across complex healthcare environments, **Enrich** augments identity records with valuable contextual data, **Verify** strengthens identity assurance across interactions, and **Match Transparency & Explainability** provides visibility into identity decisions for governance and auditability.

Together, these capabilities provide the trusted identity intelligence required for AI agents to operate accurately, consistently, and at enterprise scale.

Enabling Enterprise-Scale Agentic AI

Verato applies AI where it adds value, such as improving recommendations and operational efficiency, while maintaining reference-data-driven identity resolution as the authoritative foundation for identity decisions.

This approach enables healthcare organizations to move beyond isolated AI use cases and build agentic AI solutions that are accurate, explainable, governed, and ready to scale.

Before Scaling AI Agents, Scale MDM

Healthcare organizations are moving quickly from AI experimentation to enterprise-scale deployment. As AI agents become more autonomous, their success will depend not only on the sophistication of the models they use, but also on the quality of the identity data that guides every decision and action.

AI agents require trusted identity as the basis for accurately understanding the patients, providers, organizations, and relationships across the healthcare ecosystem. Without that foundation, even the most advanced AI systems inherit fragmented records, incomplete context, and inconsistent information.

Ultimately, successful agentic AI depends on three essential qualities of enterprise identity data:

completeness, correctness, and context. Organizations that strengthen these capabilities through enterprise MDM will be better positioned to deploy AI that is accurate, explainable, and scalable.



Verato MDM Cloud™ provides the trusted identity foundation that helps healthcare organizations move from AI experimentation to enterprise-scale agentic AI.



References

1. Gartner. Lack of AI-Ready Data Puts AI Projects at Risk.
<https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk>
2. Gartner. Data Quality: Best Practices for Accurate Insights.
<https://www.gartner.com/en/data-analytics/topics/data-quality>
3. The New England Journal of Medicine.
The Social Determinants of Health — Moving Beyond Screen-and-Refer to Intervention.
<https://www.nejm.org/doi/10.1056/NEJMms2211450>
4. Deloitte. Business and IT leaders report AI agents are scaling faster than their guardrails.
<https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html>
5. Studies in health technology... Social Determinants of Health Data Quality at Different Levels of Geographic Detail.
<https://pubmed.ncbi.nlm.nih.gov/37203650/>
6. Javelin. Identity Fraud Losses Totaled \$43 Billion...
<https://javelinstrategy.com/press-release/identity-fraud-losses-totaled-43-billion-2022-affecting-40-million-us-adults>
7. The National Institute of Standards and Technology (NIST). Trustworthy and responsible AI.
<https://www.nist.gov/trustworthy-and-responsible-ai>



Verato®, the identity intelligence company, solves the problem that drives everything else—Knowing Who Is Who™. Proven in healthcare. Trusted everywhere.

Verato MDM Cloud™ delivers unprecedented identity intelligence by combining industry-leading identity accuracy with enterprise-wide data mastering to unify, enrich, and manage identities across systems of insight, engagement, and record —unlocking more accurate, actionable data insights.

With complete, accurate, and trusted identity intelligence, Verato improves business results by powering exceptional consumer experiences and patient outcomes, fueling business growth, accelerating AI readiness, boosting productivity, and reducing costs across every business process that consumes identity. For more information, visit verato.com.