PricewaterhouseCoopers Services LLP

Enabling trusted agentic AI: The role of master data management

Enterprise AI is evolving beyond systems that help organisations generate insights towards systems capable of taking autonomous action. While predictive AI improved forecasting and generative AI enhanced creativity and decision support, both primarily informed human decisions rather than executing them. Agentic AI builds on these capabilities by enabling systems to reason, make decisions, collaborate across functions, take autonomous actions, and continuously improve through real world feedback. Figure 1 illustrates the evolution from predictive AI to generative AI and ultimately Agentic AI marking the next stage of enterprise AI, where intelligence not only informs business decisions but also drives their execution.

Market signals highlight how quickly this shift is unfolding. **Gartner** predicts that **40%** of enterprise applications will embed tasks specific AI agents **by 2026**, up from under **5%** in **2025**. **Gartner's** best case scenario projection predicts that agentic AI could drive **approximately 30%** of enterprise application software revenue **by 2035**, surpassing **$450 billion**.[1]

Figure 1: Why 'doing' is gaining popularity

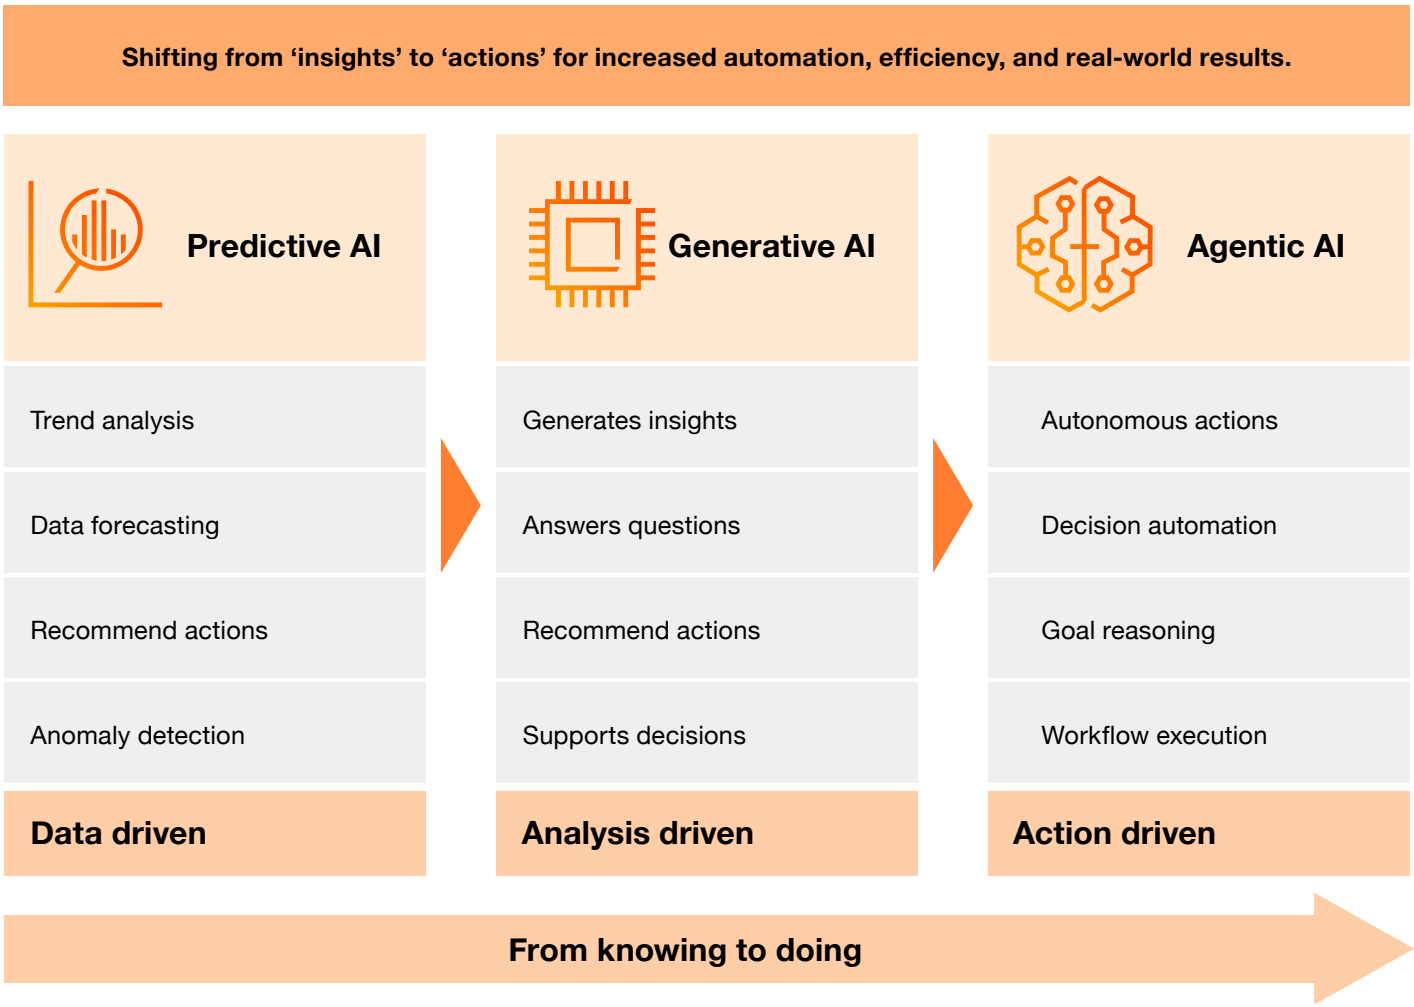

1 Gartner Press Release, Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025, 26 August 2025. https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

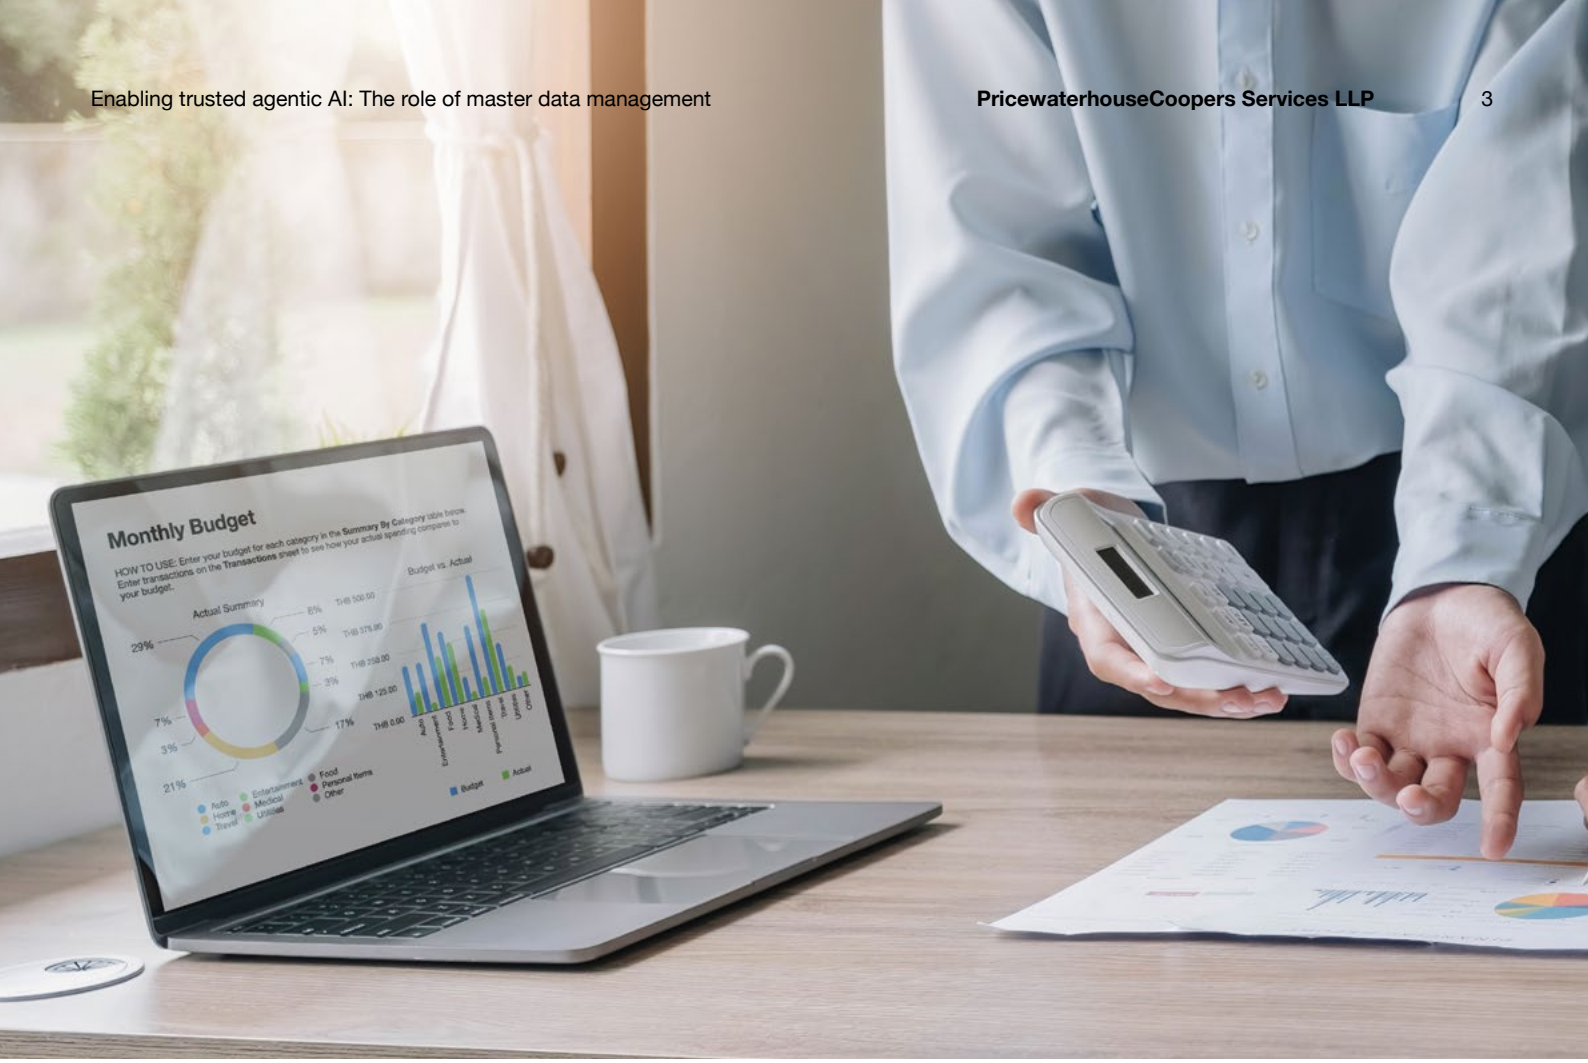

Data-related challenges in agentic AI

Enterprises are moving from AI that recommends to AI that executes. The real challenge, though, isn't the model, it's whether agents can act on enterprise data that is accurate, trusted, and governed. Data inconsistency becomes an operational risk. Taking the wrong action can be far more damaging than a wrong prediction.

The gap between AI ambition and enterprise data is already visible in how organisations are preparing themselves for adopting AI. **63%** of organisations either do not have or are unsure if they have the right data management practices for AI.[2] Poor data quality costs organisations at least **$12.9 million** a year on average, according to Gartner research from 2020.[3]

Consider a bank using agentic AI to support a personal loan journey. If the customer's records are duplicated across systems, the agent may not know the real 'who' behind the application, creating **identity ambiguity**. If income, liabilities, credit history, and relationship values are not connected, **data silos starve the agent of context** and could lead to recommending an unsafe decision. If the data is outdated or inaccurate, **poor data quality increases both cost and operational risk**, creating delays, rework, and risk. Without clear governance on which data can be trusted and used, **governance gaps stop scaling**, limiting agentic AI to controlled pilots.

For agentic AI, poor master data quickly becomes a business execution risk. When agents act on fragmented customer, product, or vendor data, they can make inconsistent decisions, trigger incorrect actions, or fail to deliver the intended business outcome. This is why trusted master data becomes essential to scaling agentic AI safely and effectively.

---

2 Gartner Press Release, Lack of AI-Ready Data Puts AI Projects at Risk, February 26, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk

3 Gartner insights, Data Quality: Best Practices for Accurate Insights, https://www.gartner.com/en/data-analytics/topics/data-quality

The master data management (MDM) advantage: Establishing an authoritative data foundation for agentic AI

Why do MDM and agentic AI belong together?

You don't need MDM to run agentic AI. The moment the agents work on anything meaningful such as making decisions about customers, products or suppliers, coordinating across systems, they are touching master data. And if the master data is unreliable, your agents will be too. For example, If the same supplier exists under multiple names across different systems in your organisations, the AI procurement agent might recommend different suppliers for same product or inflate spend analysis. Therefore, any significant agentic orchestration involving masters requires MDM as an authoritative data foundation as illustrated in Figure 2.

Figure 2: Why does agentic AI need an authoritative data core?

Agentic AI systems don't just analyse, they plan, decide, and act in multi-step loops across various tools and enterprise systems with minimal human intervention, thus elevating the role of data from 'available' to 'ready for action'.

When AI systems act autonomously:

Data quality becomes operational risk

Context gaps become execution failures

Inconsistent facts become unexplainable outcomes

Essential questions for any enterprise before deploying agentic AI:

What data can our agents safely trust?

How do we prevent autonomous actions from being based on stale or conflicting facts?

Can we explain and defend every agent decision after it happens?

What is 'Ready for Action' data?

- **Clean and unified:** Duplicates resolved, inconsistencies addressed, standardised formats.
- **Governed and compliant:** Security, privacy, consent, and policy boundaries enforced by design.
- **Context-rich:** Entities with attributes, relationships, interactions, and relevant unstructured content.
- **Instantly accessible:** Delivered with low latency where agents operate.

This is where master data management (MDM) becomes important not just as a data hygiene function but as core infrastructure for safe autonomy.

How MDM enables trusted and scalable agentic AI

Before agentic AI can act, it needs to know two things—who this is and what it's for. MDM helps in creating a single, trusted record for every entity that agentic AI depends on. It creates a persistent ID that entity follows across every system, process and agent interaction, ensuring that the agent acts from the right answer, not the nearest one.

Agents know who they are dealing with. The duplicate records are resolved and clearly defined relationships between entities such as customers to account, suppliers to product are established to remove any ambiguities and conflicts for agents to act. They know how much to trust, what data they are looking at because MDM provides data quality score which acts like a confidence signal.

MDM standardises vocabulary, product codes, units of measure and hierarchies for agents so that they can communicate in a common business language. When any data changes due to match and merge, MDM captures the lineage, audit trail and change history, thereby providing transparency and making autonomous AI auditable, trustworthy, and explainable.

Agents always work on what is valid and latest data. Through real-time integration and change-data-capture, MDM ensures that agentic AI operates on live data. For any exceptions, it can be routed to human stewards, ensuring the connection between human oversight and automation stays intact.

In an agentic environment, good enough data is often not good enough as siloed, inconsistent and outdated data would only fuel poor automation, increased risk and loss of trust. If your data is not ready, then your agents aren't either.

Figure 3: What to look for in an agentic-ready MDM platform?

Continuous ingestion and quality

Quality and freshness are maintained continuously not just at onboarding

Built-in automation

AI-augmented matching, quality scoring, and anomaly detection reduce manual effort and scale stewardship

Speed to value

Preconfigured best-practice models and mappings accelerate time to trusted data

Multidomain context

Unifies customer, product, supplier, location, and more in a single backbone for cross-process decisions.

Governance and access control

Policy-enforced access, PII controls, and audit-ready lineage

AI-ready interfaces

Low-latency APIs/events and MCP integration to plug trusted data into agent workflows

The following examples demonstrate how MDM can enable agentic AI prevent fraud and enhance care pathways.

Enterprise-grade fraud detection powered by governed, connected and resolved entity intelligence

Consider this, a loan application passes every standard check but the applicant's name exists under three different names across the systems. Two of those profiles share a device with accounts which are linked to suspicious transactions flagged in another separate system that never interacted with the system which approved the loan. By the time anyone connects the dots, the damage is already done.

Modern financial crimes and fraud are intrinsically network problems and not just rule-based ones. These network related issues become visible only when your data is connected, resolved and governed well-enough to see them as whole.

This is precisely what MDM enables—it resolves fragmented identity records into single trusted entity view, maps relationships between customers, accounts, and devices and builds a governed entity-graph. This connected foundation enables graph-based analytics which can monitor shared devices, mismatched identities, indirect links, and unusual money movements which can become visible when seen together. I

For agentic AI, this foundation allows:

1. Autonomous triage of high-risk networks
2. Context-rich investigations with explainable lineage
3. Faster interdiction with reduced operational cost.

The outcome is not only reduced instances of fraud related losses but also greater regulatory confidence and audit defensibility at scale.

How MDM can help autonomous healthcare agents

Consider this, an agent detects early signs of sepsis in post-surgical patient and fires an alert. The attending physician has just handed over to a colleague, and the care plan was updated an hour ago. The patient's consent records are sitting in separate billing systems which have not been reconciled with the clinical records.

As the context is incomplete, the alert goes to the wrong person. Nobody made a mistake as AI did exactly what it was supposed to do. The problem was the data it worked from.

An AI alert can only be useful if the agent sending it knows who the real patient is, who is responsible for their care, what their current plan is and whether the proposed plan is permitted or not. Else, AI agents can overburden clinicians with redundant or incorrect alerts, miss opportunities for intervention, or even cause automation which can be dangerous for the patient. Healthcare MDM addresses this at the foundation. It preserves one persistent, authoritative view of every patient's care context—their current care plans, assigned teams, escalation pathways, consent scopes, and clinical standards. All of this is governed, connected, and kept current across every system.

For agentic AI, this foundation enables:

- Context aware clinical execution which is aligned with care guidelines.
- Continuous patient risk surveillance across longitudinal care states in real time
- Safe and auditable automation under human supervision.

This results in a trustworthy, accountable autonomous AI in healthcare that is made possible by quicker, safer decisions with less work for clinicians, fewer missed interventions, and better results.

Way forward

Organisations need to consider MDM as a strategic enabler instead of an IT back-office project since investing in agents without strengthening master data limits scale and increases risk. They need to start by focusing on master domains with highest business impact, wire them into your agent workflows, and use data quality scores as operational guardrails. Organisations that succeed in this:

01 Reduce operational risks and hallucinations because agents make their decisions in verified, high-quality facts

02 Scale autonomous decision making safely as agents act on resolved identities, regulated semantics, and defined regulations.

Agentic AI thrives when it is based on an authoritative data core and MDM develops that foundation, entity by entity, policy by policy, and decision by decision; and ensure that businesses can move more quickly and confidently.

At PwC, we deliver end-to-end MDM services from MDM strategy and roadmap to implementation and managed operations. Contact us to learn how trusted master data can help you scale agentic AI with safety and speed.

Contact us

Prakash Suman
Executive Director, Data Governance Lead
PwC India
prakash.suman@pwc.com

Abhishek Chaurasia
Director, Data Governance Lead
PwC India
abhishek.chaurasia@pwc.com

Somoj Chowdhury
Director, MDM Lead
PwC India
somoj.chowdhury@pwc.com

Authors

Somoj Chowdhury
Abhishek Chaurasia
Annanya Swaroop
Abhisek Kundu

HS/August 2026 - M&C 54613