Informatica from Salesforce | Deloitte.

White Paper

AI-Powered Data Management

How AI agents are rewriting the rules of enterprise data management

By Patrick Taleng and Tobias Rebele

Where data & AI come to LIFE

The conundrum of data management

There is a paradox at the heart of corporate data ambitions: the more enterprises seek to use data as a strategic asset, the more likely they are to become mired in the mechanics of managing it. Data volumes are rising, regulatory expectations are multiplying and self-service demands are surging. Yet the number of skilled data stewards is not scaling up at the same rate.

Faced with this imbalance, organizations have turned to automation. Scripts have been written and workflows have been established, but most actions still depend on human triggers.

AI agents promise to release the pressure. Unlike previous generations of data tooling, they go beyond simply accelerating human action to starting to substitute for it, proposing, executing, and adapting within predefined boundaries. Enterprises benefit from faster workflows and autonomous action, with rule enforcement woven directly into the data plane. The question becomes less about how to automate and more about how to govern tools that can react to changing conditions and act without waiting for instruction.

Gartner suggests that embedding AI into data integration tools could **reduce manual intervention by 60%**[1]. Yet the solution extends beyond efficiency alone. The real shift lies in moving from reactive governance to what might be termed **controlled autonomy** – a model where intelligent agents handle routine data stewardship tasks while governance teams define the boundaries, escalation thresholds, and trust parameters.

This does not mean replacing data stewards. It means augmenting them with a new operational layer: a software-mediated workforce capable of continuously monitoring, classifying, remediating, and orchestrating data flows at machine speed.

Meet the authors



Patrick Taleng

Director, AI & Data,
Deloitte Middle East



Tobias Rebele

Data Platform Sales Leader DACH, Middle East and LATAM
Informatica

This is not about replacing data stewards. It is about augmenting them with a new operational layer

[1] Informatica is named a Leader in the 2025 Gartner® Magic Quadrant™ for Data Integration Tools for the 20th time: https://www.informatica.com/data-integration-magic-quadrant.html

Why traditional models are cracking

Data volumes and complexity have outpaced the human bandwidth available to govern them. As a result, the stress points in today's data governance models are becoming increasingly visible.

Research from Informatica[2] reveals that 67% of enterprises have been unable to transition even half of their GenAI pilots to production. For 43%, the quality, completeness and readiness of data remains primary barriers to scaling.

On one hand, new business initiatives, digital products, and ecosystem partnerships create constant requests for new datasets, attributes and quality checks. Meanwhile, the surge in generative AI and machine learning programs has created both a greater dependency on high-quality, well-understood data as well as a new source of data proliferation. Models produce new derived assets faster than governance structures can catalog them.

On the other hand, Data stewards – often embedded within business functions – are rarely dedicated full-time to stewardship responsibilities. At the same time, governance offices are rarely resourced to scale to demand. Budgets don't expand in line with data growth, and headcount increases are difficult to justify, especially when most tasks are seen as "administrative" rather than strategic.

Against this backdrop, governance processes risk becoming procedural checkpoints rather than strategic enablers. Ticket queues grow. Trust erodes. Business users route around governance to get things done.

The result? Fragmentation disguised as agility.

AI agents enter at this precise friction point.



[2] Informatica: CDO Insights 2025: https://www.informatica.com/resources.asset.3af4e34dceb6210a82be7fa135fc7a57.pdf

The rise of AI agents in the data stack

Most data management automation has focused on either workflow orchestration or static rule enforcement, moving data from A to B, applying pre-configured validations, or flagging exceptions for manual review.

AI agents introduce a step-change in capability by combining three attributes:

1. **Perception**
   Understanding the state of data systems
2. **Reasoning**
   Selecting the right remediation pattern based on context
3. **Autonomy**
   Executing actions within defined boundaries without waiting for human approval

Think of AI agents as context-aware virtual workers. Embedded within data management platforms, they operate 24/7/365 to classify newly ingested assets, suggest lineage documentation, apply policies, and surface evidence of schema drift.

The shift is subtle but meaningful. Where traditional governance focused on slowing things down for control, AI-driven models seek to accelerate action responsibly, removing human intervention not to bypass governance, but to industrialize it.

From reactive governance to proactive control

In many organizations, governance activates only after something goes wrong: a data quality violation surfaces, a regulatory audit reveals lineage gaps, or a business unit complains about inconsistent metrics. AI agents proactively scan for potential issues and resolve them before they propagate downstream.

This posture creates value in two ways:

1. **It shortens time-to-remediation.** Problems that may have gone unnoticed are detected in near real-time
2. **It creates a continuous feedback loop**, feeding every action back into the governance fabric. Over time, metadata evolves from a static documentation asset into a living control system

That evolution should start with an assistive mode, where agents generate recommendations for review. The next step is conditional execution. Ultimately, organizations can move to governed autonomy, where agents act independently but within clearly defined policy boundaries. Escalation triggers are codified, not assumed.

This is controlled autonomy in practice: software executes, governance supervises.



Impact on business data stewards

For business data stewards, the most time-consuming tasks arise from validation and monitoring. Much of their effort is spent checking fields, chasing inconsistencies, and ensuring that data behaves according to governance rules.

AI agents shoulder this burden. They patrol datasets in real time, test entries against business rules, flag anomalies and append useful contextual insights drawn from across the organization's data estate. Instead of waiting for scheduled dashboard refreshes, stewards gain a continuous stream of observations and alerts.

Discovery also becomes less of a treasure hunt. Newly arrived datasets are scanned, tagged, and slotted neatly into existing taxonomies without human prompting. Relationships between data assets are suggested rather than painstakingly mapped.

Without these constraints, business data stewards can focus on high-value tasks that machines can't credibly do: such as determining if a dataset is strategically relevant, aligning definitions across business outcomes and promoting greater adoption.

Newly arrived datasets are scanned, tagged, and slotted neatly into existing taxonomies without human prompting. Relationships between data assets are suggested rather than painstakingly mapped

Key steps to successful implementation

As AI agents start taking on more data stewardship tasks, governance operating models must evolve accordingly. Most organizations will need to establish a new layer of oversight: an autonomy governance board or equivalent body tasked with approving escalation pathways, confidence thresholds, and audit standards.

This implies a tiered decision structure:

Defining policies:

- Humans establish rules and set the thresholds that trigger action
- AI agents turn policies into executable logic and operate within those parameters

Handling exceptions:

- Humans review alerts and approve high-risk actions
- AI agents act once predefined confidence or risk thresholds are met

Active monitoring:

- Humans oversee processes using dashboards and audit trails
- AI agents continuously execute remediation, detection, and monitoring tasks

Instead of relying on manual signoffs throughout the data management process, checkpoints can become algorithmic guardrails. Metadata is tagged to indicate when an agent took action and if a human override was necessary. This creates a transparent, inspectable audit trail of governance decisions.

But caveats remain. As AI agents advance in speed and capability, they will be able to perform a wider array of tasks. But human stewards must remain accountable for outcomes. The shift is from performing actions to defining the conditions under which those actions should be performed.

Risks and constraints

Despite their promise, setting agents to power data governance raises new concerns.

Poorly trained agents might misapply policies or amplify biases already present in data sets. Over-automation could obscure systemic errors until they surface downstream. Regulatory frameworks, meanwhile, still assume human accountability, not stewardship via software.

Auditability will be essential to successful implementation. Every autonomous decision must leave a traceable record, including the rule invoked, data affected and reasoning applied. Explainability in model governance may have been a niche concern in the past, but with agents it becomes central to data governance itself.

The **economic dimension** must also be considered. While AI agents may cut labor costs, they can increase dependency on proprietary infrastructure and compute resources. Short-term efficiency gains must be weighed against long-term lock-in. As with any automation dividend, there is an accompanying maintenance cost.

Marching toward agentic data management with Informatica

No data management company has embraced this transition more effectively than Informatica. Our Intelligent Data Management Cloud (IDMC) already provides the backbone of corporate data estates around the globe, integrating, cleansing, and cataloging information across hybrid and multi-cloud architectures. Now it is being refitted around AI agents, transforming a collection of smart tools into a self-orchestrating fabric.

In 2025, Informatica was **named a Leader**[3] in the 2025 Gartner® Magic Quadrant™ for Augmented Data Quality Solutions for the 17th time. Key strengths highlighted by Gartner included CLAIRE GPT and AI Copilot. CLAIRE GPT offers a user-friendly natural language interface, while CLAIRE AI Copilot automates routine data management tasks. The integration of Cloud Data Access Management within IDMC improves data governance and security.

The CLAIRE AI engine already leads machine-learning to vital tasks like enriching metadata, detecting anomalies, and scoring data quality. The new generation of agents extends that from analysis to action.



Data-quality agents do more than flag violations; they correct them automatically based on learned approval histories and confidence thresholds.



Schema and integration agents monitor source systems, simulate downstream impacts and apply approved transformations.



Governance-workflow agents triage exceptions, prioritizing by risk and logging every decision for audit.

These agents are embedded directly into IDMC's runtime, not bolted on as advisory modules. Governance therefore occurs in motion. Each data ingestion becomes a checkpoint; each steward override becomes a training event; each metadata update refines the system's lineage graph. The platform evolves toward continuous, policy-aligned execution – a textbook example of controlled autonomy in practice.

In parallel, Informatica is bringing agents into IDMC's data marketplace. Agents will be able to maintain dynamic trust scores for each dataset, updating lineage and quality indicators continuously. Discovery and governance operate in harmony. Business users can locate the most relevant data sets and see a trust score attached.

With AI agents woven into data management, governance takes on a market dynamic; self-balancing, feedback-driven, and perpetually in motion.

[3] Informatica is named a Leader in the 2025 Gartner® Magic Quadrant™ for Augmented Data Quality Solutions for the 17th time: https://www.informatica.com/data-quality-magic-quadrant.html

The takeaway: Controlled autonomy is the new standard in governance

The introduction of AI agents into enterprise data management represents a foundational shift in how governance is conceived and executed. High-performing organizations will embrace **controlled autonomy**, a model in which intelligent agents execute repetitive stewardship tasks at scale, while governance teams set the parameters of trust, escalation, and accountability.

Integrated across the data lifecycle, agents can create a continuous fabric of compliance and assurance. They make governance less a barrier and more a property of the system itself.

But technology isn't destiny. While AI agents transform enterprise data management, competitive advantage lies in keeping humans at the core of strategic decision-making. The shift to controlled autonomy will require initial focus on high-friction, high-cost decisions, linking agentic automation to measurable outcomes, and identifying the capabilities required to balance autonomy with oversight as conditions evolve.

The real risk is no longer losing control to automation, but clinging to manual governance models while competitors build systems that learn, act, and improve – without waiting for permission.

Get in touch

To find out more, contact the experts:



Patrick Taleng

Director, AI & Data,
Deloitte Middle East
ptaleng@deloitte.com



Tobias Rebele

Data Platform Sales
Leader DACH, Middle
East and LATAM
Informatica
trebele@informatica.com

Or contact us **here**.

About Deloitte

Deloitte has been around for 175 years now, and in an ever more complex world our purpose is to make an impact that matters. To our clients, people and society. Whether it's for providing assurance, digital transformation, adaptive business, M&A advice or the great societal challenges in the fields of Energy, Food, Health, Mobility, Security or Work – at Deloitte we create solutions. Together with our clients.

Through multi-disciplinary collaborations we connect our clients with strong partners and ecosystems. We develop new ideas, insights and solutions for the most challenging and complex issues of our time. It's how we accelerate breakthrough outcomes.

Learn how Deloitte's approximately 415,000 people worldwide make an impact that matters at **www.deloitte.com.**

About Informatica

Informatica from Salesforce is a leader in AI-powered enterprise cloud data management. Its Intelligent Data Management Cloud (IDMC) platform enables organizations to connect, manage and unify AI-ready data across the enterprise. With capabilities spanning data cataloging, integration, governance, quality, privacy, metadata management and master data management, Informatica supports a broad partner ecosystem and helps customers unlock the full value of their data and AI initiatives.

About Salesforce

Salesforce is the #1 AI CRM, empowering companies to connect with their customers in a whole new way through the power of artificial intelligence, data, and trust. For more information about Salesforce (NYSE: CRM), visit: **www.salesforce.com**.

Worldwide Headquarters
2100 Seaport Blvd.
Redwood City, CA 94063, USA
Phone: 650.385.5000
Fax: 650.385.5500
Toll-free in the US: 1.800.653.3871

informatica.com
linkedin.com/informatica
x.com/Informatica

CONTACT US

Where data & AI come to



IN09-5214-0326

© Copyright Informatica LLC 2026. Informatica and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and other countries. A current list of Informatica trademarks is available on the web at **https://www.informatica.com/trademarks.html**. Other company and product names may be trade names or trademarks of their respective owners. The information in this documentation is subject to change without notice and provided "AS IS" without warranty of any kind, express or implied.