



NUCLEUS
RESEARCH

The value of Informatica's headless agentic integration platform

ANALYST

Alexander Wurm

The Bottom Line

Enterprises are moving AI agents into production, but an agent is only as reliable as the data it can reach. Most of that data sits behind fragmented ETL and ELT pipelines, and moving, cleaning, and transforming it consumes the largest share of data-team time. Informatica recognizes this expanding use and has delivered multiple new releases intended to extend the ease of developing and managing integrations. Notably, Informatica made its Intelligent Data Management Cloud headless, exposing integration as governed services. Every capability from ingestion, transformation, governance, quality, master data management and CLAIRE agents, will become available for agents to call from any surface through the Model Context Protocol. This is expected to deliver 27 to 58 percent improved data engineer productivity and allow Informatica customers to reclaim the capacity previously lost to hand-built integration work.

Overview

The integration layer is where enterprise AI stalls. Data teams spend 40 to 80 percent of their time connecting, cleaning, and transforming data before anyone can use it. Every new source or target adds more code to maintain. The move from ETL to ELT pushed transformation into the warehouse, yet it multiplied the pipelines a team has to govern. Agentic AI raises the stakes. An agent acting on stale or ungoverned data returns confident, wrong answers, and the cost of that error grows as agents take real actions on their own.

Informatica addressed this challenge with its release of headless data management. The approach exposes IDMC integration, ingestion, quality, and governance as reusable services that agents and applications call on demand, rather than functions locked inside one interface. It also extended ingestion and change data capture to Microsoft Fabric and added connectors for Databricks and Snowflake, keeping data movement open across clouds after the Salesforce close.

Informatica made its Intelligent Data Management Cloud headless, exposing integration as governed services.

Informatica

Headless Integration

Headless data management exposes each capability, from ingestion and transformation to quality and lineage, as a service an AI agent can call through the Model Context Protocol. The company kept these services open across Amazon Web Services, Google Cloud, Microsoft, Databricks, and Snowflake. That decision matters for integration buyers. Rather than route data into a single vendor's environment, teams feed agents from the platforms they already run. CLAIRE, Informatica's intelligence layer, now operates across those same environments to automate discovery, mapping, and transformation, so the choice of cloud no longer dictates the choice of tool.

A new Data 360 connector adds real-time, bidirectional data flow with end-to-end lineage, so teams trace every field to its source.

Agentic Integration and Expanded Connectivity

Informatica's agentic integration capability, generates integration pipelines rather than requiring hand-coded mappings, extending the automation CLAIRE already applies to discovery, mapping, and transformation. A new Data 360 connector and scanner adds real-time, bidirectional flow with end-to-end lineage, so teams can trace where each field originated. For Databricks, a named connector for Lakebase ingests and transforms operational data, reaching transactional sources. For Snowflake, Informatica added governance for open Iceberg tables and propagation of row-level access policies, so data

stays controlled as it moves into the lakehouse. Informatica also announced agents for data quality and master data management, though the integration services carry the clearest near-term value for teams building agent-ready pipelines.

Expected benefits

Headless, agentic integration is expected to deliver gains across data engineer productivity, data-team capacity, and the trustworthiness of the data agents act on.

Improved Data Engineer Productivity

Integration has long depended on specialists writing and maintaining pipeline code. Headless services and AI-assisted pipeline generation shift that work from coding to direction, letting engineers assemble and adjust pipelines without rebuilding them. Nucleus expects data engineer productivity to improve 27 to 58 percent as AI-assisted integration replaces manual work, with new pipelines reaching deployment up to 90 percent faster. It also lowers the skill barrier, letting less specialized staff build pipelines that once required a dedicated engineer. A fixed team can then support more sources, more targets, and more agent use cases without adding headcount.

Reclaimed Data-Team Capacity

Data teams lose 40 to 80 percent of their time to connecting, cleaning, and transforming data, which leaves analysis and AI enablement underserved. Automating ingestion, change data capture, and transformation returns that time to higher-value work. Nucleus expects teams to redirect capacity from pipeline maintenance toward data products and agent governance, the work that separates a working AI program from a stalled one.

Trusted Data for AI Agents

An agent is only as good as the data behind it, and ungoverned pipelines feed it stale or unverified records. Exposing integration, quality, and lineage as governed services means the data an agent reads stays current and traceable to its source. Nucleus expects adopters to cut the rework and failed outputs that follow from bad inputs, and to move agents into production with fewer data-quality holds. Governed integration turns the pipeline itself into the control point for trustworthy AI.

Any agent can consume Informatica's integration, ingestion, and quality services as governed building blocks, with no pipeline rebuilt.

Nucleus expects AI-assisted integration to improve data engineer productivity 27 to 58 percent and cut deployment time up to 90 percent.

Looking ahead

Nucleus expects the integration layer to become the control point for enterprise AI. As agents move from pilots to production, the question shifts from which model to use toward whether the data feeding the agent is current, governed, and traceable. Headless integration and open, multi-cloud access position that layer as the trusted source agents draw on.

For buyers, the guidance divides by readiness. Teams running legacy PowerCenter should evaluate migration to the cloud platform now, a move Nucleus has tied to lower ongoing cost. Organizations building multi-cloud agents should evaluate the headless services already available, where governed integration is the immediate gap. Teams without an agent roadmap can wait for agentic integration, and adopt once agentic pipeline generation is proven in their stack.

Informatica has made agentic integration, which generates data pipelines without hand-coding, a priority on its near-term roadmap.