CAPCO

Agentic Data Management

Transforming data operations for the future

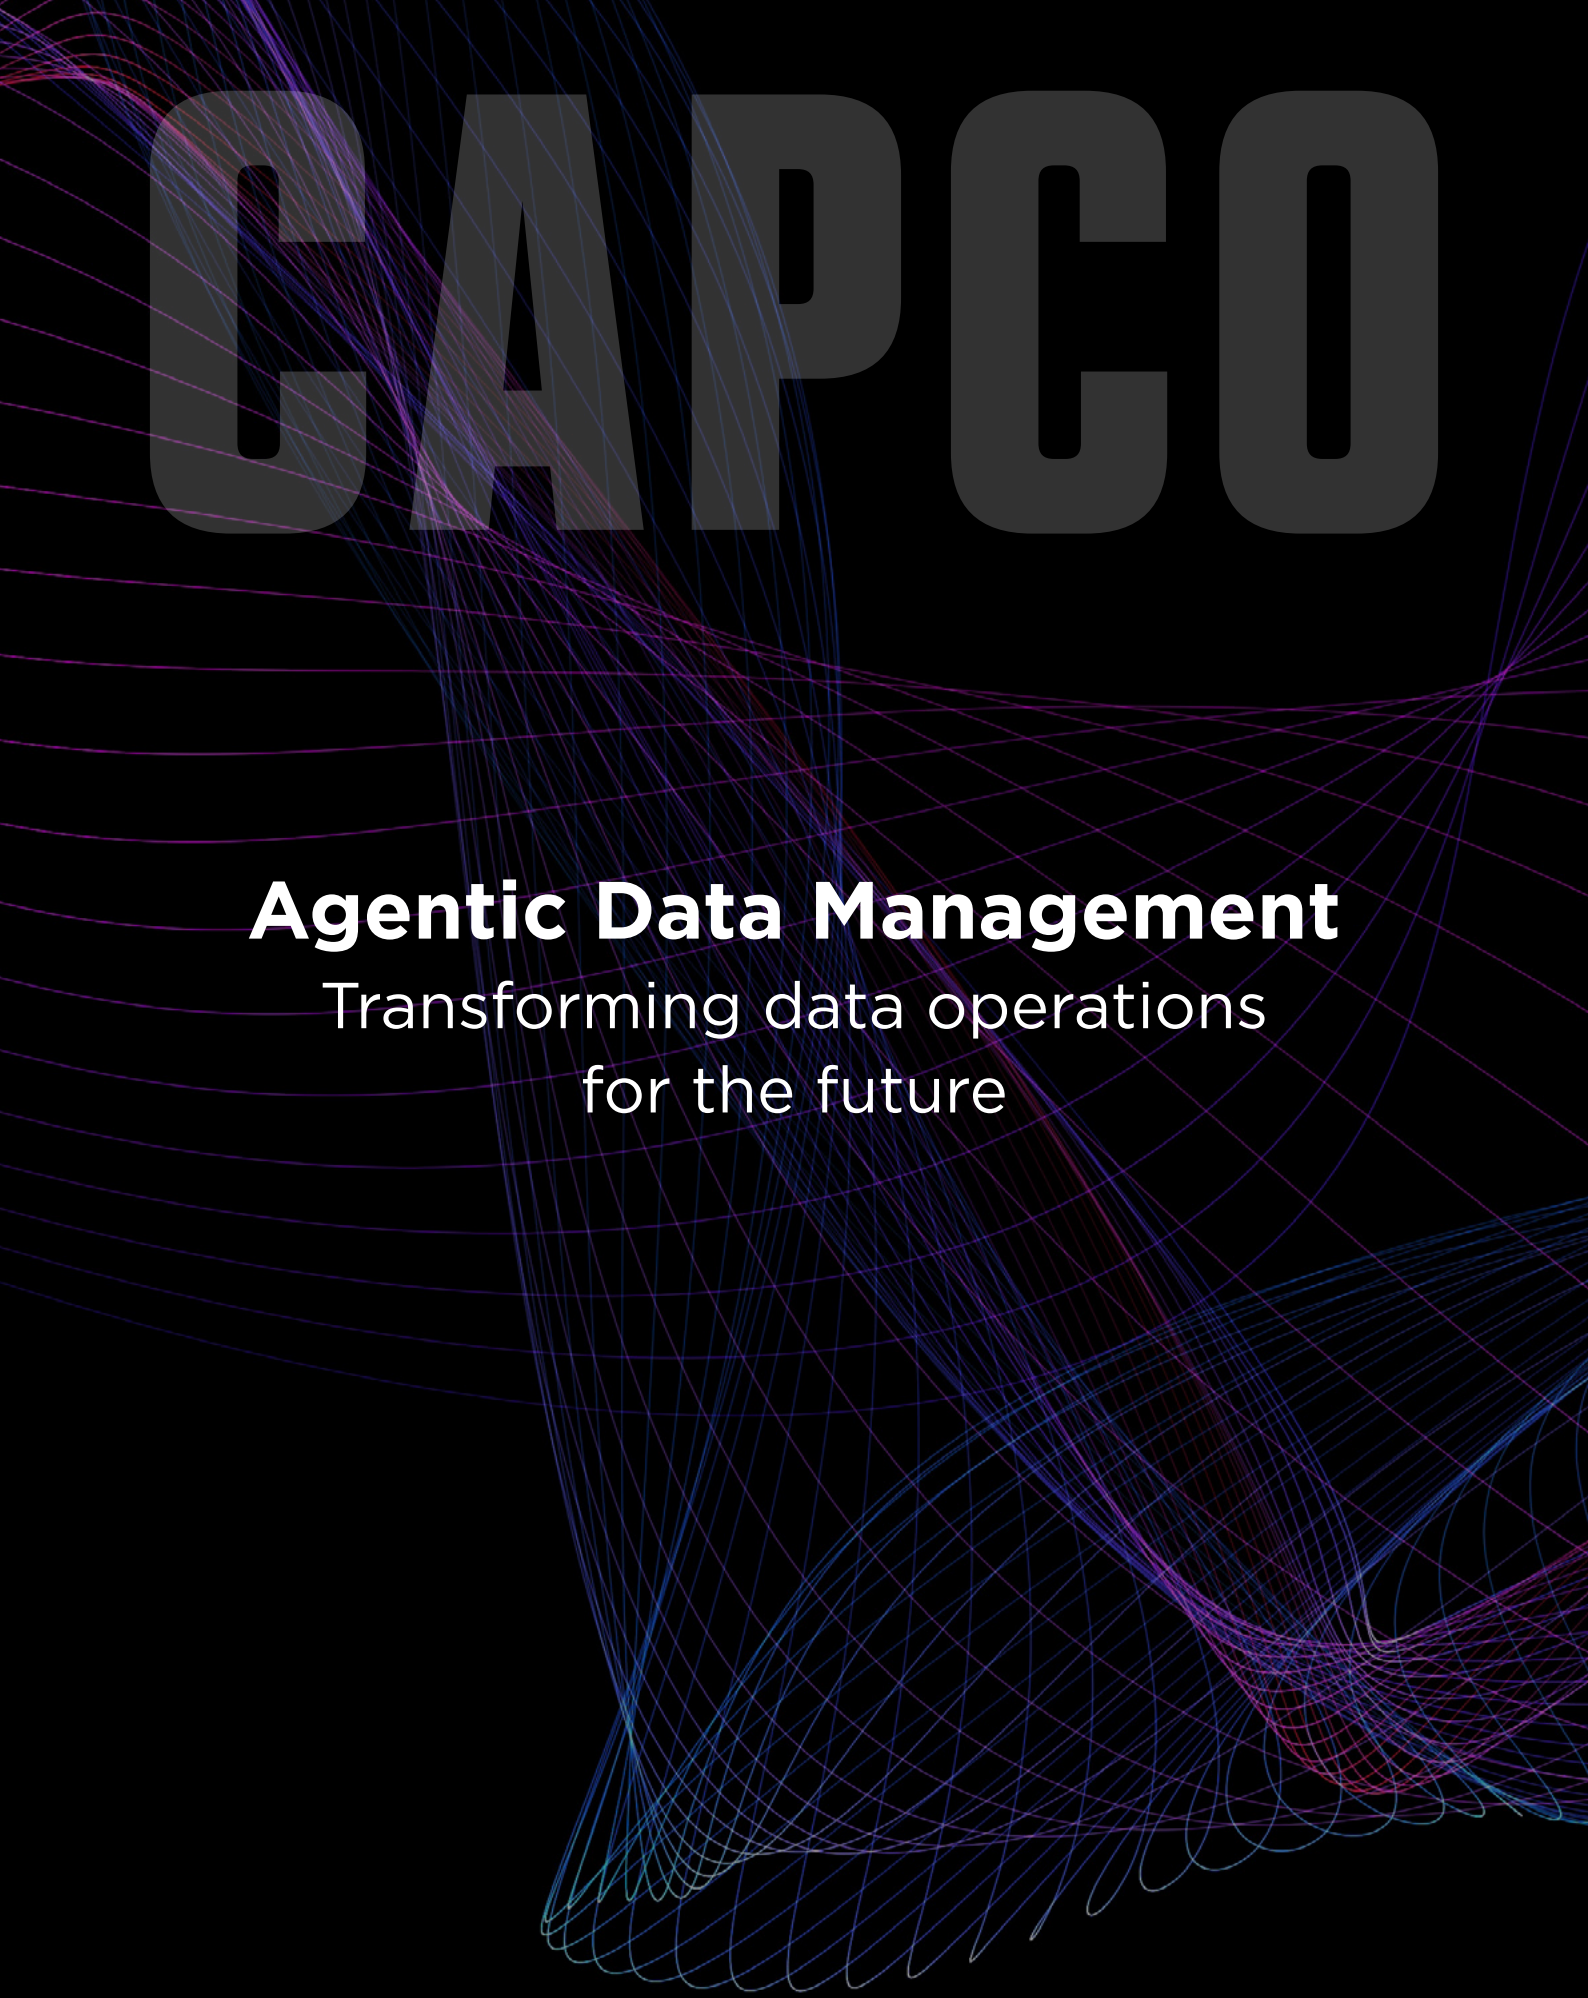

a **wipro** company

Traditional data management has largely depended on manual, people-driven methods which have been sufficient in demonstrating regulatory compliance but ineffective in materially improving data quality, consistency or usability. Chief Data Offices (CDOs) face mounting pressure to deliver tangible ROI, accelerate innovation and reduce costs while ensuring compliance and governance. Agentic data management offers a path forward.

The perceived cost-effectiveness challenge facing CDOs reflects a legacy focus on compliance and regulation-driven data management. As data management advances from siloed automation to intelligent, self-improving systems, CDOs need to reposition data management as an enabler of data usability and value creation – delivering real-time insights, driving agent autonomy in data management and ensuring data is fit for purpose so it can be confidently used across the organization.

Most firms currently have dedicated data management tools (data catalogs, DQ tools, lineage solutions etc.) in place, but automation within these is often siloed, designed as point solutions and requiring human integration, configuration, rule writing and oversight. Some are making additional progress via augmented, Machine Learning-enhanced metadata scanning, lineage extraction and DQ checks.

Agentic data management uses AI agents to support and manage data activities across the lifecycle – planning, reasoning, executing, learning and continuous improvement – while keeping humans in the loop for oversight and complex decision-making. Routine and well-defined tasks can be handled with minimal supervision, while more nuanced scenarios rely on human judgment, validation and governance.

This approach shifts data operations from manual, fragmented processes to intelligent, goal-driven systems that guide users through complex workflows, reduce effort and promote consistency and control.

As organizational knowledge matures and processes become more established, the level of automation can progressively increase, enabling faster, more informed decision-making at scale without compromising accountability.

With reference to the maturity scale in Figure 1 below, organizations typically sit somewhere between 2 and 4, with ambitions to realize the full value of automation and agentic solutions as they work toward their ‘North Star’.

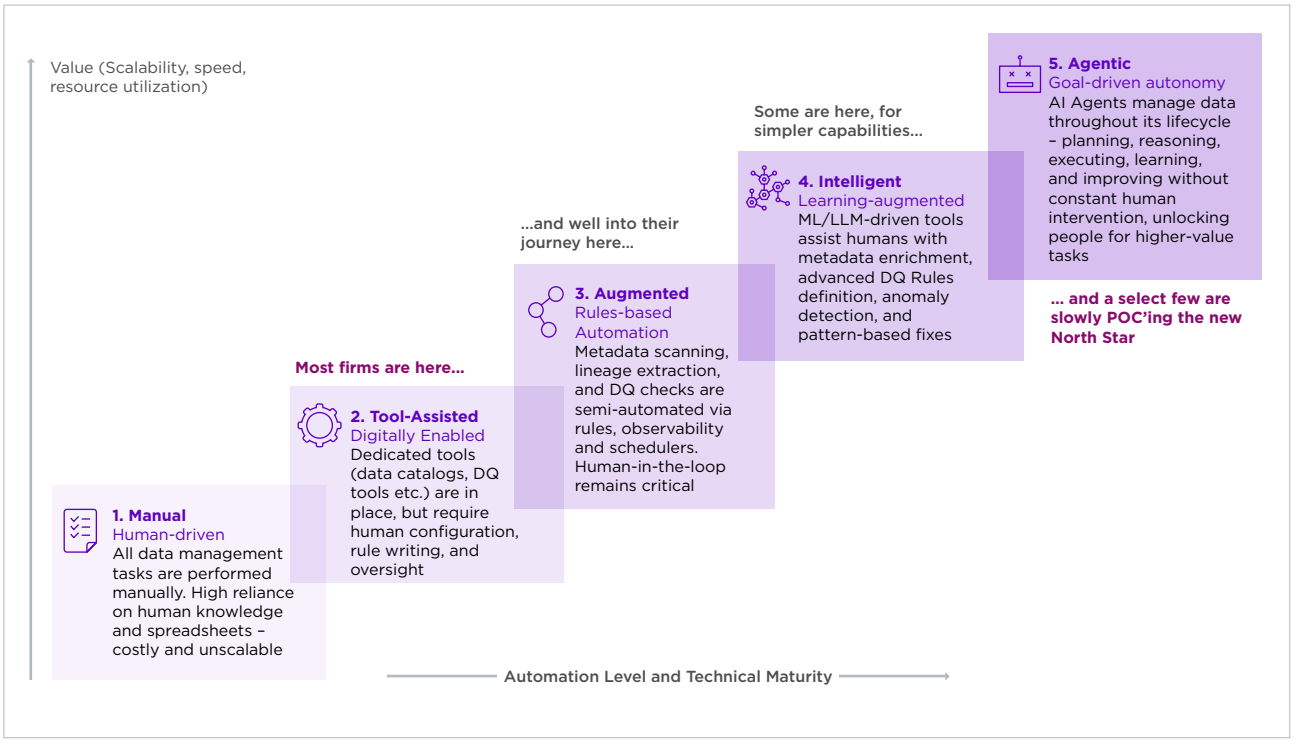

Figure 1: The path to agentic data management

Agentic data management: an opportunity for CDOs

The strategic importance of data is no longer in question; and it will remain a mission-critical asset. Firms operate in ecosystems where decisions must be made in real-time, based on ever-expanding volumes of structured and unstructured data. Yet, many organizations still rely on traditional, human-led approaches to the end-to-end management of data, ranging from metadata and lineage to data quality and remediation. Such approaches are costly, fragmented, and incapable of scaling with demand.

Data volumes and complexity continue to grow, and manual approaches cannot keep pace. Business demand is growing faster than human capacity. At the same time, skilled data professionals are stretched thin, leaving organizations struggling with talent gaps.

Meanwhile, ongoing vigilance is key to ensuring compliance, but this is a burden that human-led processes alone are struggling to sustain.

Agentic data management is an emerging approach whereby autonomous AI agents manage data across its lifecycle – from discovery and classification to lineage, quality control, and resolution. These agents are both decision-makers and doers, capable of planning, reasoning, executing and learning with humans remaining in the loop for governance, accountability and high-impact decisions. In practice, these agents with humans-in-the-loop function like a team of digital counterparts that can take on end-to-end data management tasks, ensuring scale, accuracy and speed.

Core characteristics of agentic data management include:

- **autonomy** – AI agents operate independently to fulfill objectives
- **goal-driven execution** – humans set high-level targets (e.g. "ensure 99% KYC completeness") that agents deliver against
- **collaboration** – agents coordinate with each other to complete complex workflows, with human oversight for complex or risky decision-making
- **continuous learning** – performance improves through feedback and outcomes
- **embedded guardrails** – ensuring compliance, security, and governance are maintained throughout.

The result is a continuously improving operating model where AI agents and human teams work in close partnership, combining automation with expert judgment to deliver reliable outcomes.

The mechanics of agentic data management

Agentic data management operates through a layered architecture where humans and AI agents work in tandem, each focusing on their strengths. At the top is the **human interaction console**, which provides a simple interface for humans to set goals, monitor activity, and step in when needed to approve or override agent decisions. These high-level objectives are then translated into action through a suite of **policy, orchestration, and monitoring agents**.

For example, a policy agent converts human goals into machine-readable constraints and success criteria, while an observability agent ensures agent actions remain compliant with regulatory and risk standards. A dynamic workflow agent breaks down policies into smaller executable tasks, and an orchestration agent coordinates execution and inter-agent collaboration across the workflow. Access and usage rules can be set, and humans are in the loop for oversight. This ensures organizations benefit without losing control.

Once policies are translated into tasks, a modular set of specialized **data management agents** handle execution. These include agents for metadata discovery, lineage mapping, data quality monitoring, root cause analysis, resolution, and escalation. Each agent performs its role autonomously, but in coordination with others, ensuring end-to-end automation of data management processes.

Supporting this ecosystem is a **knowledge base** which is a shared memory layer that includes data catalogs, management tools, process flows, policies, and system logs. This serves as the collective intelligence of the platform, enabling agents to learn, adapt, and refine their actions over time.

Together, this structure transforms data management from a manual, fragmented effort into a coordinated, autonomous system. Humans define strategic intent, policies provide direction, agents execute with speed and precision, and the knowledge base drives continuous learning ensuring that data operations remain scalable, compliant, and responsive to business needs in real time.

Our blueprint architecture shows a visual representation of how humans, intelligent agents and a shared knowledge base operate as a unified, self-orchestrating ecosystem (see Figure 2).

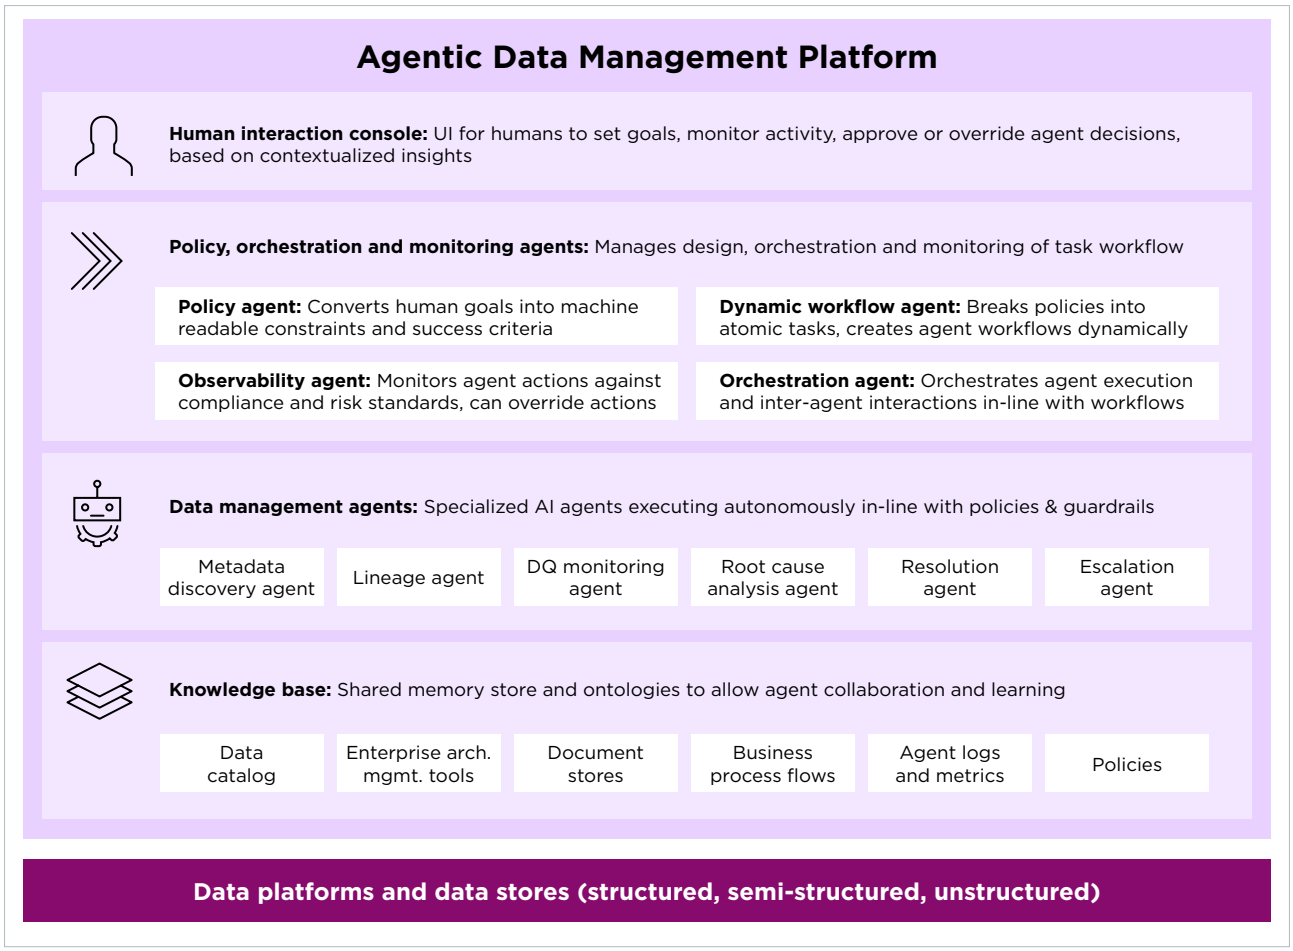

Figure 2: Agentic data management platforms

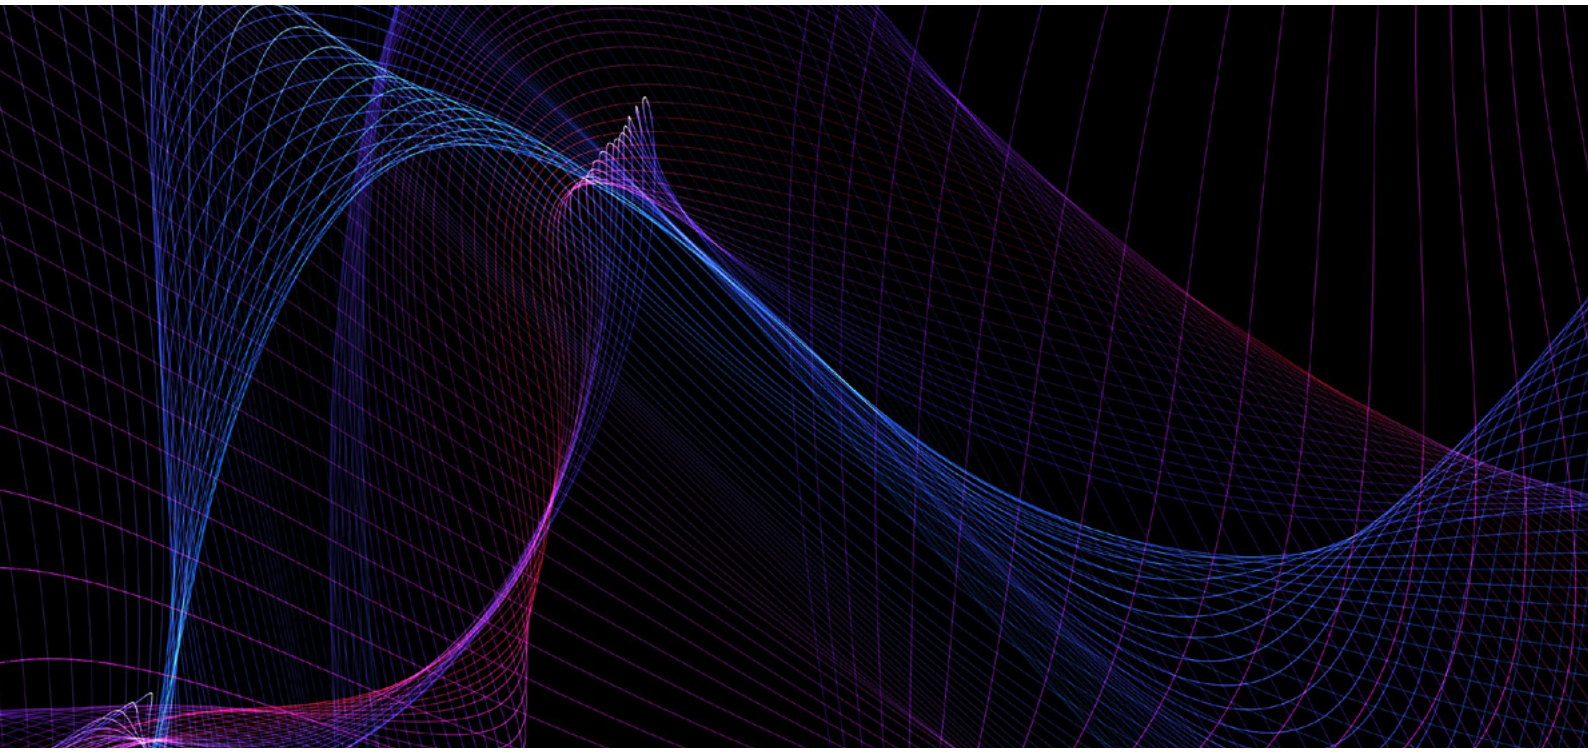

The benefits of agentic data management

Agentic AI offers a strategic opportunity to automate, optimize, and future-proof data operations, creating the foundation for smarter decisions, stronger trust, and scalable governance. It transforms the way organizations handle complexity by embedding intelligence directly into the data fabric and enabling CDOs to move beyond operational work toward strategic leadership. As Figure 3 highlights, the key benefits are:

- user empowerment
- smart metadata management
- proactive and real-time data quality management
- process and decision automation
- a future-ready foundation
- enhanced CDO effectiveness.

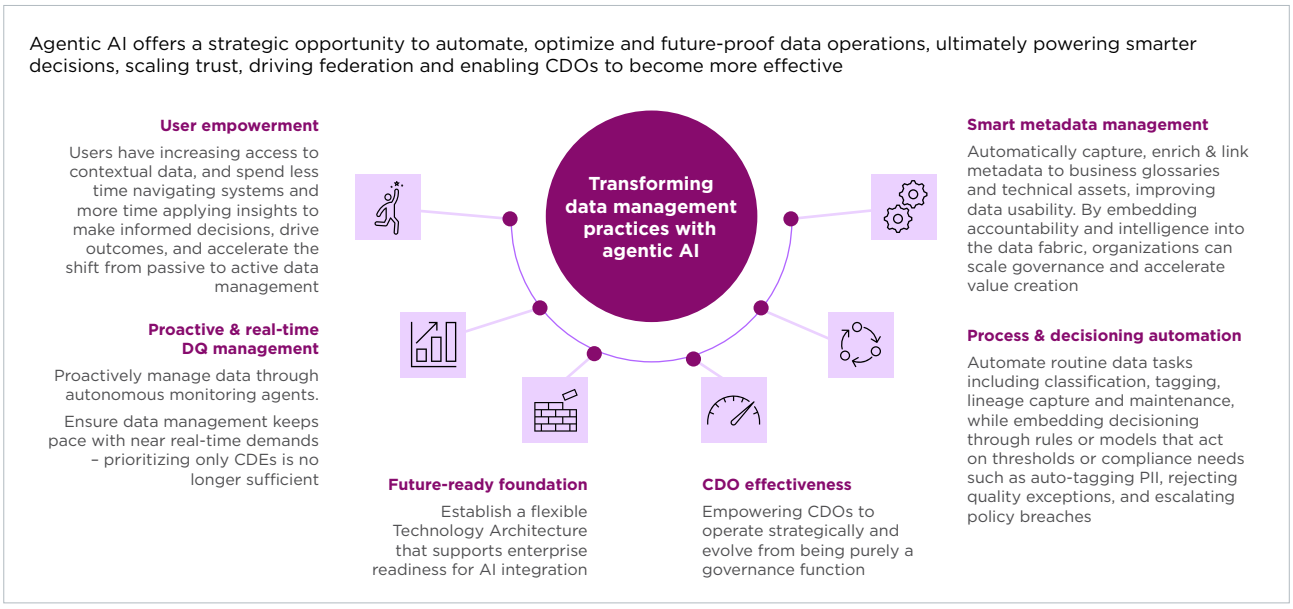

Figure 3: Benefits of agentic data management

Collectively, these benefits position agentic data management as an operational upgrade, along with being a strategic lever for competitiveness, efficiency and trust.

With agentic AI, coordinated agents carry out the six phases of the data management lifecycle at scale – automating discovery, lineage, monitoring, diagnosis and resolution, while humans set objectives and guardrails and step in only for oversight and escalation (see Figure 4).

Phase	Traditional human-led approach	Agentic w/ human-in-the-loop
Discover Discover, classify and catalog data assets	Data stewards manually inventory new sources; discovery is reactive and siloed. Classification relies on business knowledge and inconsistent tagging	**Metadata discovery agent** continuously scans sources, catalogs new assets, and auto-tags using ontologies. Semantic classification is updated in real time
Trace Map lineage flows	Lineage is manually documented or semi-automated through code reviews, ETL mappings, or SME interviews; often incomplete and outdated	**Lineage agent** automatically parses code, data feeds, data stores, and APIs to maintain a live, system-wide lineage map. Detects impact from upstream changes
Monitor Continuously assess data quality	DQ rules are defined manually or via basic profiling; dashboards monitored reactively; anomalies often missed or mis-prioritized	**DQ monitoring agent** auto-generates rules based on data semantics and history. Continuous validation via rules and anomaly detection, surfacing meaningful alerts
Diagnose Identify root causes of issues	Issue triage requires manual inspection of logs, identification of patterns, analysis lineage maps, and pipeline traces. Root cause often unclear	**Root cause analysis agent** correlates exception patterns with past issues and control weakness to pinpoint root causes quickly and accurately
Suggest Recommend and simulate fixes	Engineers write SQL fixes or adjust pipelines manually where relevant. More complex root causes require longer for solutioning	**Resolution agent** recommends fixes (e.g., fallback data, schema correction, control uplift, sourcing change), simulates impact, and explains resolution confidence
Fix Resolve or escalate issues	Manual escalation through ticket queues or email chains; often lacks complete context or audit trail	**Escalation agent** packages full lineage, risk impact, and fix options; routes issues to governance or ops leads based on severity thresholds, for human to approve fix

Figure 4: phases of the data management lifecycle

Real world applications

Banking & Payments:
Single Euro Payments Area (SEPA)
In banking and payments, agentic data management can transform how institutions handle complex, high-volume transactions such as those within the Single Euro Payments Area (SEPA). Traditionally, fragmented data, intricate lineage across payment systems, and manual remediation create inefficiencies and risks.

With Agentic AI, specialized agents autonomously discover and classify payment data, trace flows across systems, and continuously monitor quality. When anomalies occur, root cause analysis agents quickly identify issues, while resolution and escalation agents recommend or execute fixes in real time.

Capital Markets & Wealth Management: securities pricing
In capital markets, securities pricing is highly sensitive to fragmented data feeds, inconsistent vendor inputs, and complex lineage across multiple systems. Traditional methods make it difficult to pinpoint errors quickly, creating risks for valuation accuracy and compliance.

Agentic data management addresses this by enabling agents to autonomously scan and catalog inbound pricing data, trace lineage across instruments and vendors, and continuously monitor for anomalies.

In wealth management, the same approach ensures smarter client data management, allowing firms to optimize service quality, deliver personalized insights, and maintain data integrity across portfolios.

Insurance: managing claims
Agentic AI can transform insurance claims management by autonomously handling fragmented and inconsistent data across systems, continuously monitoring data quality, tracing claim lineage, and diagnosing root causes of errors.

By recommending or even executing fixes such as resolving duplicate claims, correcting claim classification codes, or escalating incomplete cases, Agentic AI helps insurers streamline assessments, reduce delays, minimize compliance risks, and deliver faster, more accurate outcomes for policyholders.

Energy: energy grid load forecasting
In the energy sector, Agentic AI enables more transparent and resilient grid load forecasting by discovering and cataloging fragmented inputs, tracing lineage across models, and continuously assessing data quality.

By identifying root causes of anomalies, Agentic AI helps tackle inconsistent reporting, undocumented flows, and transient anomalies to strengthen forecasting accuracy and operational reliability.

Mapping a path forward

Recognizing that not all organizations are ready for full agentic adoption, we can help institutions to establish a North Star and clear interim states, ensuring a clear vision that guides near-term decisions and ensures those decisions align with their longer-term goals.

We recommend running a focused Proof of Concept (PoC) to prove value quickly, then expanding into a broader end-to-end model once the benefits are clear and aligned to the organization's goals.

- **Agentic data management framework:** Define the vision, guiding principles, target operating model, solution blueprint, and a prioritized roadmap to establish the foundation for Agentic AI adoption
- **Human-agent operating journey:** Design how humans and agents collaborate, mapping user roles, journeys, and oversight flows to ensure agents augment productivity while maintaining human control and accountability.
- **Agentic architecture design:** Define the target technical architecture, integration approach, technology selection, and governance guardrails to support scalable and compliant deployment
- **Proof-of-Concept (POC):** Identify high-impact use cases, develop pilot implementations, and validate measurable value to demonstrate the potential of Agentic AI.
- **Scaled implementation:** Extend the proven concepts through platform engineering, product and UX design, and automated orchestration
- **Change and workforce enablement:** Prepare the organization for adoption through impact assessments, role and skill transformation, training, and enablement.

Throughout we leverage our accelerators – including our AI Labs, reference ontologies, workflow playbooks, intelligent data management framework, tooling UX design and reference architectures – and an outcome-based delivery mindset to derisk adoption, secure early wins, and embed compliance and governance from day one.

CDOs must prepare the data workforce of the future – and the next era of data management will be defined by the collaboration between humans and agents. Agents will handle scale and complexity, while humans define goals, set guardrails and ensure ethical use. Together, this partnership will define the future of data operations.

At Capco, we focus on enabling scale and reducing cost from strategy to execution, fast-tracking agentic adoption, and delivering reliable outcomes. With our frameworks, accelerators and expertise, enterprises can reimagine their data operations to drive efficiency, compliance, and business value.

Agentic data management is not just an efficiency play, but a strategic enabler – one that allows organizations to innovate with confidence and govern with control. To find out how we can support your transformation journey, please get in touch with our experts.

Authors

Furnisha Lycett
Principal Consultant, Data
Furnisha.Lycett@capco.com

Helen Needham
Managing Principal, Data
Helen.Needham@capco.com

Joseph Forooghian
Managing Principal, Data
Joseph.Forooghian@capco.com

About Capco

Capco, a Wipro company, is a global management and technology consultancy redefining transformation across the financial services and energy industries. Capco leverages the power of AI and our deep domain expertise to help our clients move faster, make smarter decisions, and drive greater impact. Our award-winning Be Yourself at Work culture and diverse talent drive bold, forward-thinking ideas and lasting change.

To learn more, visit www.capco.com or follow us on LinkedIn, Instagram, Facebook, and YouTube.

Worldwide Offices

APAC
Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurugram
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST
Dubai

EUROPE
Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto

SOUTH AMERICA
Rio de Janeiro
São Paulo

capco.com

CAPCO
a wipro company

JN_7469