Modern Data 101 Community

October 2026

Why AI-Native Transformation Stalls at the Tool Layer

Why 74% of AI's economic value goes to just 20% of organisations

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Table Of Contents

1. Executive Summary
2. From Digitisation to AI-Native Process Logic
3. The Tool-Layer Trap
4. What the AI-Native Enabling Layer Must Provide
5. The Enabling Layer Model
6. The Enabling Layer Test
7. The Organisational Operating Model is Part of the Architecture
8. A Practical Roadmap for AI-Native Transformation
9. What an AI-native Enterprise Should Ultimately Achieve
10. Conclusion
11. References

© 2026 Modern Data 101 (by TMDC). All rights reserved.

1

Executive Summary

From The Modern Data 101 Community

This report highlights AI-native transformation: moving AI from isolated tools to governed, end-to-end business processes through an enabling layer of data products, shared context, feedback loops, and decision-level governance

Most large organisations now use AI in at least one function, yet few have changed how their core processes are designed, governed, and improved. Deloitte's 2026 State of AI in the Enterprise survey finds that 37% of organisations use AI at a surface level, with little or no change to their underlying processes. PwC's 2026 AI Performance Study finds that 74% of AI's economic value is captured by 20% of organisations. Lysenko (2026) describes AI-native as the level of digital transformation where AI becomes part of end-to-end process logic.

This volume argues that the step from AI tools to that level is blocked by a missing enabling layer: governed data products, a shared semantic context, closed feedback loops, orchestration and decision-level governance that connect the AI investments an organisation already has. It draws on insights from the major consultancies and other leaders to locate where the layer is weakest, sets out its components and requirements, and offers a diagnostic that tests whether a given process has moved from tool use into a closed-loop cycle.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

2

From Digitisation to AI-Native Process Logic

Digital transformation has moved through recognisable stages. Digitisation moved paper processes into software. Digital optimisation automated individual steps with workflow and integration tools. AI augmentation added models to existing steps, such as a forecast inside a planning tool or a classifier in a ticket queue. Lysenko's framework adds a further level. At the AI-native level, AI is embedded in the whole process logic, spanning strategy, management, operations, control, and improvement, and those elements connect into a single cycle.

The distinction matters because each earlier stage improves parts of a process while leaving its design intact. Deloitte's AI Institute draws the same line between adopting AI and transforming with it. Adoption deploys tools across an unchanged process, while transformation redesigns the process itself. Its 2026 AI Pulse Check reports that only 12% report the most mature state of AI, where it can run end to end and humans audit outcomes rather than approve each step. It also recommends auditing for process change, because adoption metrics and transformation metrics measure different things.

Deloitte's enterprise AI operating model analysis adds that AI-native is not simply the next phase of digital transformation. Earlier digital programmes modernised how a business ran its existing operating model. An AI-native enterprise makes cognition part of that operating model, so that decisions, actions, and learning run through a common layer.

The commercial stakes of this distinction are now measurable. PwC's 2026 AI Performance Study surveyed 1,217 senior executives across 25 sectors. It finds that growth from industry convergence is the strongest factor in AI-driven financial performance, ahead of efficiency gains alone. For this volume, AI-native redesign is an economic question as much as a technical one.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

3

The Tool-Layer Trap

Tools stall transformation for three connected reasons.

First, tools optimise tasks, and task gains are easy to report. A drafting assistant or a forecasting model can show measurable improvement in isolation, so organisations declare progress that never reaches the process. Deloitte's survey finds that 30% of organisations are redesigning key processes around AI, and a further 37% report using AI at a surface level, with little or no change to underlying business processes. Cisco's explanation of AI-native systems offers a practical test: remove the AI and ask whether the product's main value disappears. Applied to a process, the question becomes whether the process would run any differently without AI.

McKinsey's 2025 State of AI survey shows the same divide from the performance side. Only 39% of respondents attribute any EBIT impact to AI. High performers, about 6% of respondents, are nearly three times as likely as others to fundamentally redesign their workflows.

Only 39% of respondents attribute any EBIT impact to AI. High performers, about 6% of respondents, are nearly three times as likely as others to fundamentally redesign their workflows.

Second, tools do not share context. SAP's systems of context analysis describes an AI feature inside one application that can summarise an invoice but cannot see that the same supplier has an open service dispute, a delayed shipment, and a renegotiated contract elsewhere. Deloitte's cognitive layer paper describes the same symptom as pilots that do not compound and agents that act without context.

Third, the data underneath is not ready. The central obstacle is activation: data that is discoverable, contextual, and trustworthy at the moment a decision is made. 68% of respondents say their data is not clean or reliable enough for AI use cases. Only 30% use AI for automated data preparation. The report argues that AI exposes the activation gap rather than closing it.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Organisational readiness adds a fourth dimension. McKinsey's three horizons research finds that only 13% of leaders in the enablement horizon report meaningful enterprise value from AI. That share rises to 24% in the automation horizon, and only 11% of leaders say they have reached the reinvention horizon. McKinsey attributes 48% of the difference between value capturers and others to organisational readiness, compared with 25% for personal readiness.

Only 13% of leaders in the enablement horizon report meaningful enterprise value from AI.

These causes share a root. Each tool was bought to improve a step, and nothing was built to connect the steps.

4

What the AI-Native Enabling Layer Must Provide

These requirements are not parallel features. They form a dependency chain, and the enabling layer is best understood as a set of contracts at the points where one process level hands work to the next.

It is more like multiple requirements, one chain. Shared context makes a signal interpretable. A signal that is interpretable but arrives late cannot support a timely decision. A decision that is made but not recorded cannot be learned from. Each requirement is the precondition for the next, so a failure at one point does not stay local. It becomes bad input for the link that follows.

This is why organisations that buy strong models and monitoring tools still see little change: the chain breaks at a handoff, and the tools only operate on either side of it.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Handoffs are where transformation breaks. In a traditional process, a handoff is a person passing a document, and ambiguity gets resolved because the receiver can ask a question. In an AI-native process, the receiver is often a model or an agent that reads the output of the previous step without asking. Ambiguity that a human would clarify becomes a silent error. This problem can be described at the data level, through semantic ambiguity, temporal misalignment, and format inconsistency. We think the same failure modes appear at the process level, where the unit of failure is a decision rather than a record.

Orphaned automation is the characteristic failure

As per EY, "Agentic AI adoption is creating new governance challenges, with 26% of respondents whose organzation uses agentic AI admitting that their organization cannot detect unauthorized AI agents operating internally."

The broader pattern is automation that acts without a named owner for its outcomes or its exceptions. We call this orphaned automation. It is more dangerous than manual error because it runs continuously and its errors compound before anyone notices. A practical test follows from this: for every automated step in a process, can the organisation name the person accountable for its results and the person who handles its exceptions? Where the answer is no, the enabling layer is incomplete, however capable the technology.

Agentic AI adoption is creating new governance challenges, with 26% of respondents whose organization uses agentic AI admitting that their organization cannot detect unauthorized AI agents operating internally.

AI-ready to Action-ready

For an AI-native process, that standard is too narrow. The data must also be ready for action. A stock figure of 40 units tells a model very little. A stock figure that carries its reorder threshold, its supplier lead time, and the approval needed to place an order tells the process what it is allowed to do next. Our proposal is that every data product consumed by an automated step should carry the permitted action alongside the fact. This is the point where data governance and process governance become the same design problem.

Design rule

Every handoff between process levels needs a contract that specifies four things: what the data means, how fresh it must be, what action it permits, and who owns the outcome. Organisations that write these contracts before they buy tools will find that the tools fit. Organisations that buy tools first will spend the next two years writing the contracts retroactively, one incident at a time.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

5

The Enabling Layer Model

The enabling layer has five components. Each one addresses a specific failure in the tool-layer trap.

1. **Data products with contracts.** Each process consumes defined data products with a named owner, quality thresholds, freshness commitments, and embedded business rules. A product with a contract can be discovered, trusted, and reused, which removes the repeated rebuilds that keep AI projects slow.

2. **A semantic context layer.** This is the shared vocabulary, relationship model, and provenance record that lets every process interpret the same entity the same way. It is the component that makes shared context possible across systems that were never designed to talk to each other.

3. **Process and agent orchestration.** This component decides where AI sits in each process: which steps run autonomously, which propose a decision for human confirmation, and which remain fully human. Cisco's principle of automation with limits applies directly. Oversight should match the potential harm of a decision, and that boundary should be a deliberate policy choice.

4. **The feedback and improvement loop.** Process mining, KPIs, operational logs, customer feedback, and audit trails feed back into the data products and the process design. Lysenko's advantages table describes this as moving improvement from a periodic project to a permanent mechanism.

5. **Governance and decision records.** Each automated decision records its inputs, the rule or model used, the approver where one exists, and the outcome. Lineage, explainability, and memory governance sit here. This component answers the audit and accountability questions that otherwise stall AI-native programmes at scale.

One of our analyses talks about where AI sits relative to the data platform and distinguishes three models: embedded within the platform, shared as a horizontal layer, and external.

For an enabling layer that must serve many process teams, the shared horizontal model fits best. It gives teams common context and feedback without rebuilding governance for each use case.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

6

The Enabling Layer Test

The test below is a diagnostic for any single process. It turns the concept of AI-native into five questions a team can answer with evidence.

1. Context: Do all steps in the process read the same definition of the entities they handle?
2. Signal: Does the process detect changes in state in time to act on them?
3. Decision: Is each decision recorded with its inputs, the rule or model behind it, and its approver where one exists?
4. Action: Are the boundaries of autonomous action defined, with an escalation path when they are exceeded?
5. Learning: Do outcomes flow back into the data products and the process design?

Score each link as absent, partial, or connected. A process qualifies as AI-native only when all five links are connected within one governed cycle. A process can have a strong model and still fail the test if one link is missing, because the model then operates as a tool inside an unchanged process.

7

The Organisational Operating Model is Part of the Architecture

The enabling layer cannot be assembled function by function. Claims cannot build its own version of knowing, fraud cannot build its own version of learning, and underwriting cannot build a consistent picture for decisions. Each team builds differently, none shares, and the integrating value disappears. The layer has to be built once, centrally, and applied to the processes that matter most.

McKinsey's three horizons research finds that 84% of leaders in the enablement horizon and 68% in the automation horizon say their organisations are not ready for the people and culture shifts an agentic future requires. The people side is also where Accenture's research on talent is most direct.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Why existing structures constrain AI

Enterprise workflows often cross departments with separate budgets, targets, technology stacks, and decision rights. A process that appears unified to a customer may be divided internally between sales, finance, operations, service, and compliance.

AI can accelerate individual activities in this process while leaving the underlying coordination problem intact.

If every department optimises its own metrics, an agent operating within one function may make decisions that undermine the wider business outcome. A sales agent may prioritise conversion without accounting for fulfilment constraints. A service agent may optimise resolution time without considering customer retention. A finance agent may minimise short-term expenditure while increasing operational risk.

These problems are not solved by giving each agent more autonomy. They require an explicit definition of the outcome being optimised and the constraints within which the workflow must operate.

McKinsey's research on AI-driven organisational reinvention emphasises outcome-oriented work, redesigned workflows, and clearer coordination between humans and agents. Its operating-model research also indicates that the value of AI depends on deliberate organisational choices rather than a single universally applicable structure.

18% of organisations generate measurable AI value, and attributes their advantage to redefining how work gets done.

Organise around outcomes

An enterprise does not necessarily need to abandon its existing departments. It does need clear ownership of the end-to-end processes through which business value is created.

For an AI-enabled order-to-cash workflow, for example, accountability should extend across order acceptance, fulfilment, invoicing, collections, and customer experience. The relevant teams must agree on shared objectives, decision boundaries, and escalation procedures.

AI systems can then be assigned specific responsibilities within that workflow, with access to the required context and authority to perform permitted tasks.

Human roles also change. Employees increasingly need to supervise exceptions, define policies, evaluate outcomes, and make decisions where judgment remains essential. Process owners become responsible for the performance of the combined human-AI system, not simply the performance of an individual tool.

This is a material organisational change. It affects roles, incentives, governance, skills, and how work is measured.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Build reusable capabilities instead of isolated projects

A successful AI implementation should create assets that subsequent initiatives can reuse. These may include governed data products, business definitions, reusable integrations, evaluation datasets, workflow components, access policies, and observability mechanisms.

Without reuse, every new use case begins with another round of data discovery, integration, policy implementation, and operational testing. Costs and complexity grow alongside adoption.

A shared platform approach can reduce this duplication while allowing domain teams to build applications closer to their business needs. The platform should provide common capabilities without forcing every workflow into an identical implementation.

The balance matters: excessive centralisation can slow experimentation, while uncontrolled decentralisation reproduces the tool-layer problem.

An effective operating model establishes common enterprise guardrails and reusable foundations while leaving appropriate implementation decisions with teams closest to the business outcome.

A shared platform approach can reduce this duplication while allowing domain teams to build applications closer to their business needs. The platform should provide common capabilities without forcing every workflow into an identical implementation.

Change management must accompany system design

Employees need to understand how responsibilities change, when AI outputs should be trusted, how to challenge incorrect recommendations, and where accountability remains human.

Training in prompt writing alone will not prepare an organisation for this shift.

Teams need practical experience with redesigned workflows, escalation paths, evaluation methods, and the limits of automated decision-making. Leadership must also align incentives with the intended outcomes so that employees are not rewarded for maintaining processes that the transformation is meant to improve.

AI-native transformation is consequently a technology, operating-model, and change-management program. Treating any one of these as a downstream activity creates avoidable friction.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

8

A Practical Roadmap for AI-Native Transformation

AI-native transformation should progress through measurable changes to business operations and the architecture that supports them. Enterprises do not need to rebuild their entire technology estate before starting, but they must address the constraints that prevent isolated successes from scaling.
The following plan connects business priorities, data foundations, workflow redesign, governance, and operating-model change.

Phase 1. Select the Business Outcome

Establish the business baseline before selecting the technology. Identify a workflow where AI could materially improve cost, speed, quality, revenue, or customer experience. Map its current process, decision points, data dependencies, failure costs, and ownership.
Exit criterion: A measurable business objective, a named process owner, and an agreed baseline.

Phase 2. Assess the Data and Context

Determine whether the information available can support the intended decisions. Assess data quality, freshness, access, lineage, semantic consistency, business rules, and decision history. Identify the data products and shared context capabilities that need to be created or improved.
Exit criterion: Critical data dependencies, context gaps, and ownership responsibilities are documented.

Phase 3. Redesign the Workflow

Define how people, AI systems, and enterprise applications will work together. Separate tasks requiring AI reasoning from deterministic operations. Establish where agents can act, which permissions they need, when human approval is required, and how exceptions and failures will be handled.
Exit criterion: A tested workflow with defined responsibilities, controls, and success measures.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Phase 4. Validate Production Readiness

Evaluate business performance and operational reliability together. Test normal and edge cases, changing inputs, policy violations, system failures, and human overrides. Measure outcomes against the baseline alongside accuracy, latency, cost, security, and reliability.
Exit criterion: Evidence that the workflow meets its acceptance criteria and can be operated reliably within defined risk limits.

Phase 5. Scale Through Reuse and Feedback

Convert successful implementations into repeatable enterprise capabilities. Reuse validated data products, integrations, policies, and evaluation assets. Monitor business outcomes, investigate failures, refresh context, and extend the approach to adjacent workflows where the expected value justifies the cost and risk.
Exit criterion: Demonstrable operational value, reusable capabilities, and a sustained process for improvement.

Start With the Workflow

The first decision is which business outcome needs to improve, not which model to deploy. Strong candidates have a clearly defined problem, access to relevant data, an accountable owner who can change the process, measurable results, and risks that the organisation can govern.
The workflow should determine the technical approach. Some use cases may require retrieval over governed documents; others may need reusable data products, event-driven integration, semantic models, or agents with narrowly scoped execution permissions. Starting with these requirements reduces the risk of building a broad AI platform before validating the business need.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Establish Production Criteria Before Scaling

A successful demonstration shows that a capability may be useful; it does not establish that it can operate reliably at enterprise scale. Before expansion, teams should define acceptance criteria for business performance, contextual reliability, failure recovery, security, permissions, auditability, operating cost, human accountability, and reuse across workflows.

These thresholds must reflect the consequences of failure. An internal knowledge assistant and an agent authorised to modify financial records require different levels of assurance, oversight, and autonomy.

Build the Platform Around Demonstrated Needs

Enterprise capabilities should emerge from validated requirements rather than a predetermined preference for a particular technology stack. Data products, semantic layers, knowledge graphs, orchestration frameworks, model gateways, and governance services can all contribute, depending on the use case.

The objective is a coherent foundation that prevents teams from independently rebuilding data access, business context, policy enforcement, and observability for every AI application. This makes the data platform a strategic component of AI-native transformation: not simply a place to store information or run workloads, but a way to make enterprise information and capabilities consistently available to the workflows that depend on them

© 2026 Modern Data 101 (by TMDC). All rights reserved.

9

What an AI-native Enterprise Should Ultimately Achieve

The purpose of AI-native transformation is not maximum automation. It is a more capable enterprise: one that can use information to make better decisions, coordinate work across boundaries, adapt to changing conditions, and improve the way it creates value.
Several characteristics distinguish this direction.

Business context becomes reusable. New applications can draw on established data products, definitions, relationships, and policies rather than repeatedly reconstructing enterprise knowledge.
Workflows become coordinated. AI systems can operate across connected processes within explicit permissions, while people retain responsibility for judgment and decisions that require human accountability.
Governance becomes operational. Controls extend from data access into reasoning, execution, evaluation, and recovery. Accountability is designed into the workflow.
Architecture supports change. The enterprise can evolve models, applications, data sources, and business processes without rebuilding every dependent AI implementation from scratch.
Learning becomes part of operations. The organisation uses evidence from actual outcomes, exceptions, and failures to improve both its AI systems and the processes around them.

These characteristics do not require every enterprise to adopt fully autonomous agents or a completely new technology stack. They require deliberate architectural and organisational choices about where AI belongs, what it can do, and how the business will govern and benefit from its capabilities.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

10

Conclusion

Enterprise AI has made it easier to introduce intelligence into individual activities. The harder task is turning those capabilities into a reliable, coordinated part of how a business operates.
That task cannot be completed at the application layer alone.

Models require access to dependable data. Data needs consistent business meaning. Context needs to remain current and governed. Agents need to coordinate with operational systems. Actions need explicit permissions and accountable owners. Outcomes need to be measured and fed back into the system. The organisation must also redesign the workflows, incentives, and responsibilities that determine whether the technology creates value.

Digital transformation established much of the infrastructure on which this next stage depends. AI-native transformation extends that work by rethinking the relationship between information, reasoning, execution, and organisational learning.

The most important strategic decision is therefore not how many AI tools an enterprise should deploy. It is which business outcomes deserve redesign, which shared capabilities will enable that redesign, and how successful implementations will become reusable parts of the enterprise.

The transition from AI adoption to AI-native operations begins when organisations stop evaluating AI as a collection of applications and start designing for the outcomes those applications must collectively deliver.

© 2026 Modern Data 101 (by TMDC). All rights reserved.

11

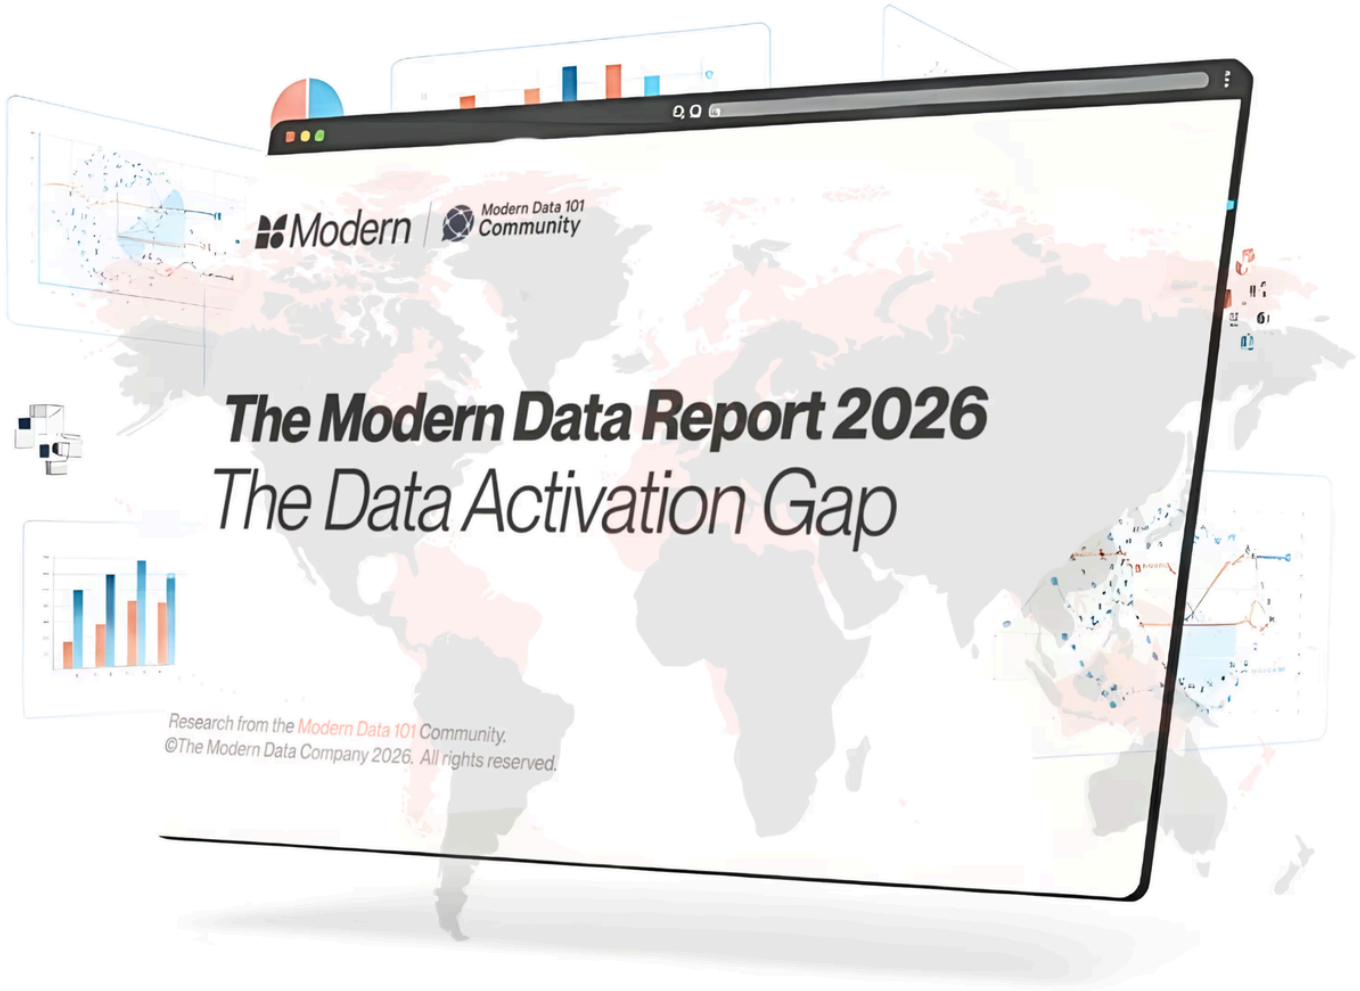
References

1. Accenture. (2026, March). Talent Reinventors report. Coverage: https://www.businessforum.ro/industry/20260323/accenture-only-18-of-firms-generate-value-from-ai-despite-rising-investment-3018
2. Calvanese, D., et al. (2026). Agentic Business Process Management: A Research Manifesto. arXiv. https://arxiv.org/abs/2603.18916
3. Cisco. What is AI native? https://www.cisco.com/site/us/en/learn/topics/artificial-intelligence/ai-native.html
4. Deloitte AI Institute. (2026). AI transformation predictions 2026 (Pulse Check series). https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/blogs/pulse-check-series-latest-ai-developments/ai-transformation-predictions-2026.html
5. Deloitte. (2026). The State of AI in the Enterprise [Press release]. https://www.deloitte.com/us/en/about/press-room/state-of-ai-report-2026.html
6. Deloitte. Sharma, P., and Limbad, B. The Enterprise AI Operating Model for AI-Native Businesses. https://www.deloitte.com/us/en/services/consulting/articles/enterprise-ai-operating-model-for-ai-native-businesses.html
7. Ericsson. (2023). Defining AI native: A key enabler for advanced intelligent telecom networks. White paper. https://www.ericsson.com/en/reports-and-papers/white-papers/ai-native
8. EY. (2026, September 15). EY survey finds that autonomous AI implementation outpaces oversight, yielding an AI governance gap [Press release]. https://www.ey.com/en_us/newsroom/2026/09/ey-survey-finds-that-autonomous-ai-implementation-outpaces-oversight-yielding-an-ai-governance-gap
9. KPMG International. (2026). Global AI Pulse Q1 2026: Executive summary. https://assets-launch.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/04/global-ai-pulse-exec-summary.pdf.coredownload.inline.pdf
10. Kumar, A. (2026). The AI-Native's Woes of Persistent Memory. Modern Data 101. The AI-Native's Woes of Persistent Memory
11. Lysenko, S. (2026). AI-native as an integrative level of digital transformation in business process management. Actual Problems of Economic Sciences. https://doi.org/10.5281/zenodo.20764061
12. McKinsey & Company. Panda, C., and Soller, H. (2026, July 29). AI-native development: Rethinking software from the ground up. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/ai-native-development-rethinking-software-from-the-ground-up

© 2026 Modern Data 101 (by TMDC). All rights reserved.

13. McKinsey & Company. (2025). The state of AI in 2025: Agents, innovation, and transformation. https://www.mckinsey.com/our-insights/the-state-of-ai-2025
14. McKinsey & Company. (2026). From adoption to impact: Three horizons of AI transformation. McKinsey Quarterly. https://www.mckinsey.com/capabilities/people-and-organization/our-insights/from-adoption-to-impact-three-horizons-of-ai-transformation
15. Modern Data Company. (2026, February 4). New Report from The Modern Data Company Finds the Data Activation Gap Is Holding Back Enterprise AI and Decision-Making [Press release]. https://www.accessnewswire.com/newsroom/en/computers-technology-and-internet/new-report-from-the-modern-data-company-finds-the-data-activation-1133420
16. PwC. (2026, January 19). PwC 2026 Global CEO Survey [Press release]. https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-global-ceo-survey.html
17. PwC. (2026, April 13). PwC's 2026 AI Performance Study [Press release]. https://pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-ai-performance-study.html
18. SAP Learning. Understanding the Need for AI-Native Approach for Enterprise Business Systems. https://learning.sap.com/courses/implementing-ai-native-enterprise-to-achieve-business-outcomes/understanding-the-need-for-ai-native-approach-for-enterprise-business-systems_cc4f12ca-112e-47a2-a2a7-422b24c4cee1
19. Scale AI. (2026, April 16). AI-Native Data Layer: Making Enterprise Data Agent-Ready. https://scale.com/blog/ai-native-data-layer

© 2026 Modern Data 101 (by TMDC). All rights reserved.



540+ Respondents, 64 Countries, 29 Industries

Access the report now

© 2026 Modern Data 101 (by TMDC). All rights reserved.

Modern Data 101 Community



Subscribe

moderndata101.com