



POINT OF VIEW

Only **10%** of Enterprises Scale **AI** Successfully

CEO's Case for Process-First, Low-Code,
Agentic Orchestration



Executive Summary

Boards do not need a bigger AI agenda. They need a sharper investment thesis. Spending has outpaced strategy, and we need to bring discipline to AI investment.

The stakes are now visible in the market. In January 2026, **BCG** reported that companies expect to double AI spending this year to roughly 1.7% of revenue, while 72% of CEOs said they are now the main decision-maker on AI and 90% expect AI agents to produce measurable returns in 2026. Yet **Deloitte** found in the same month that only 25% of organizations had moved 40% or more of their AI pilots into production. **McKinsey** added a sharper warning in April: nearly two-thirds of enterprises have experimented with agents, but fewer than 10% have scaled them to tangible value, with data limitations cited by eight in ten companies as the barrier.

These signals bear a message for the board: don't fund AI as a technology program. The winners in the next phase will not be those that accumulate the most pilots or deploy the most agents.

They will be those that:

- **Fund AI** where it can change end-to-end business outcomes, not where it only adds incremental productivity
- **Insist on data**, workflow and governance readiness before approving scaled autonomy
- **Treat orchestration and control** as board-level concerns, because unmanaged agents create unmanaged risk

With this, enterprise velocity, as it is understood, evolves into an ability to move from signal to decision to action to measurable outcome quickly and with governance. Low-code provides the speed and structure to digitize and redesign work. Agentic AI extends those flows with reasoning and action. Orchestration is what turns both into an operating model that the enterprise can trust.

For CEOs and boards, the argument comes down to three decisions.

- Whether to continue funding AI as a collection of use cases, or to reshape it around process-centric operating models
- Whether to view governance as a compliance overlay, or to recognise orchestration as the control plane of the agentic enterprise
- Whether to measure success through activity and pilots, or through cycle time, decision quality, compliance and business impact

This paper answers these questions.



Why Low-Code & Agentic AI Matter Together

Why does this matter now? Because AI has moved from an innovation topic to a governance and capital-allocation topic. BCG's January 2026 research shows that AI is now being led from the top of the house. Deloitte's January findings show that production maturity still lags ambition. And McKinsey's April work shows that scale breaks where data, architecture and process discipline are weakest. For boards, that combination should change the way AI programs are approved and governed.

The practical case for convergence is strongest in enterprise processes where delay, inconsistency or weak controls have direct economic consequences.

Consider claims handling, where faster triage and exception routing affect loss ratios and service costs. Or procurement approvals, where cycle time and policy adherence affect working capital and spend leakage. Even employee onboarding, service ticket resolution or contract review, where the real prize is not another app, but a better operating rhythm across functions. Low-code structures these flows quickly, while Agentic AI reasons across them. But value is only realized when both sit inside a process that the enterprise can observe and govern.

Low-code digitizes work. Agentic AI drives decisions. Orchestration delivers outcomes.

For CEOs and boards, the argument comes down to three decisions.

- From isolated AI tools to goal-directed execution across entire workflows
- From technology investment decisions to operating model and risk decisions
- From pilot metrics to economic outcomes such as cycle time, cost-to-serve, compliance and revenue impact

Process First is the Real Differentiator

The biggest mistake in the market is deploying agents into broken, fragmented or poorly governed processes and expecting intelligence to compensate for weak execution design. Agentic systems do not fix operational disorder; they only amplify it.

McKinsey's January and April 2026 research reinforces this point. Companies are shifting budget aggressively into AI, often at the expense of ERP, data and architecture capabilities that agents depend on. McKinsey also found that nearly two-thirds of enterprises have experimented with agents, but **fewer than 10% have scaled them to tangible value**, with eight in ten citing data limitations as a barrier.

This is the reality where:

- Agents become inconsistent because they inherit broken hand-offs, unclear exception logic and poor context
- Automation becomes fragile because systems of record, approvals and controls are not designed to support autonomy
- ROI becomes accidental because outcomes are not tied to cycle time, quality, compliance or business value

Process discipline and data readiness, therefore, sit at the center of enterprise-scale agent deployment. When enterprises lead with process, the picture changes.

- Agents can act as reliable digital workers because roles, decisions, boundaries and escalation paths are explicit
- Low-code becomes an execution layer for governed forms, workflows, approvals and human-in-the-loop interactions
- Automation becomes repeatable because enterprises can reuse process patterns, controls and orchestration logic across functions

Our Take

If the process is unstable, AI will only help the enterprise fail faster. Process-first design is what turns autonomy into measurable value rather than faster dysfunction.

This is why discovery matters. Process intelligence identifies where autonomy creates value, where human judgement must remain, and where exceptions, hand-offs and approvals need to be redesigned. It must be paired with data readiness: access controls, lineage, context integrity, retrieval quality and reliable system connectivity. Without these foundations, agents cannot perform safely or consistently at scale.

Enterprises should also be explicit about where not to apply agentic AI. Poor candidates include highly unstable processes, low-frequency edge cases, decisions with unclear accountability, and workflows where data quality, permissions or policy boundaries are unresolved. The best starting points remain high-volume processes with rules plus exceptions, measurable outcomes and clear opportunities for human-in-the-loop control.

Orchestration is the Control Lane

The most important shift in enterprise AI is the emergence of orchestration as a production requirement. The 2026 market is making this visible. At Google Cloud Next, the center of gravity moved towards agent platforms, observability and governance. At IBM Think, the narrative was even clearer: enterprises are not pulling ahead by deploying more AI, but by redesigning how the business operates with governance, data and orchestration working together.

In practice, orchestration governs the hard parts of enterprise autonomy, with:

- Agent hand-offs, sequencing and task delegation across systems, teams and digital workers
- Escalation logic, approval checkpoints, fallback rules and SLA management when work cannot proceed autonomously
- Observability, policy enforcement, audit trails, identity controls and exception routing across the full process

This is the shift from isolated automation to coordinated execution.

In a process-centric model, work is not simply stitched together through integrations. It is managed end to end with clear responsibilities for agents, applications and humans, and with explicit controls for risk, cost and compliance

Low-code matters here because it provides the governed execution surfaces that enterprise autonomy still needs, ensuring:

- Execution logic is visible and manageable via visual workflows and event-driven process models
- Consequential decisions are taken inside governed boundaries, reinforced by approval paths, exception handling and user interfaces
- Reusable patterns for human-AI collaboration across functions such as procurement, service operations, finance and HR

In other words, agents should reason and act within boundaries designed for trust. That means explainability where needed, human accountability for consequential decisions, and enough observability for enterprises to understand not only what happened, but why it happened and whether it should happen again.



Board-Level Playbook for Scaling Agentic AI

Start with process intelligence, not agent design

Map the work before automating. Inventory core processes, existing low-code assets, automation components, data sources and risk constraints. Prioritize processes where cycle time, exception rates, quality and business impact can be measured

The right question is not ‘Where can we deploy an agent?’ but ‘Where can autonomy improve an outcome we can prove?’

Build the data, workflow and governance fabric together

Combine low-code for workflow and experience design, automation where deterministic execution is required, AI agents for reasoning and action, and orchestration for observability and governance. Architect identity, permissions, lineage, auditability, prompt and policy controls, and lifecycle management as part of one system.

Scaled autonomy needs a fabric, not a set of disconnected tools.

Define control boundaries before autonomy

Apply an impact-versus-feasibility lens and redesign the process from as-is to to-be with explicit human-in-the-loop checkpoints. Decide what not to agentify.

This is the difference between controlled autonomy and unmanaged experimentation.

Set boundaries for decisions, escalation, approvals, data access and model behavior before agents go live.

Scale through reusable patterns and outcome discipline

Start with measurable pilots but design for reuse from the start. Scale through templates, reference architectures, policy guardrails and centers of enablement. Measure what matters: cycle time, resolution quality, compliance, revenue impact, cost-to-serve and employee productivity.

Enterprises scale what they can govern and prove.

Why Partner with Persistent

The future enterprise will be defined by how well it governs and coordinates work across humans, applications and agents with speed, trust and discipline.

CEOs and board members need to take this call: fund AI where it can redesign how the business operates, demand evidence that data and controls are ready for scale, and insist that orchestration and accountability are built in from the start.

Persistent comes in as an enabler of process-first AI transformation, combining low-code engineering, automation, orchestration, data and governance into an approach that helps enterprises move from isolated pilots to governed scale.

To know how we can scale AI adoption meaningfully.

[Learn More](#)

About Persistent

Persistent Systems (BSE: 533179 and NSE: PERSISTENT) is a global services and solutions company delivering AI-led, platform-driven Digital Engineering and Enterprise Modernization to businesses across industries. With over 27,500 employees located in 21 countries, the Company is committed to innovation and client success. Persistent offers a comprehensive suite of services, including software engineering, product development, data and analytics, CX transformation, cloud computing and intelligent automation. The Company is part of the MSCI India Index and is included in key indices of the National Stock Exchange of India, including the Nifty Midcap 50, Nifty IT and Nifty MidCap Liquid 15, as well as several on the BSE such as the S&P BSE 100 and S&P BSE SENSEX Next 50. Persistent is also a constituent of the Dow Jones Best-in-Class World Index. The Company has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. Persistent has also been named one of America's Greatest Workplaces for Inclusion & Diversity 2025 by Newsweek and Plant A Insights Group. As a participant of the United Nations Global Compact, the Company is committed to aligning strategies and operations with universal principles on human rights, labor, environment and anti-corruption, as well as take actions that advance societal goals. With 468% growth in brand value since 2020, Persistent is the fastest-growing IT services brand in 'Brand Finance India 100' 2025 Report.

USA

Persistent Systems, Inc.
2055 Laurelwood Road, Suite 210
Santa Clara, CA 95054
Tel: +1 (408) 216 7010
Fax: +1 (408) 451 9177
Email: info@persistent.com

India

Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 0000
Fax: +91 (20) 6703 0008

