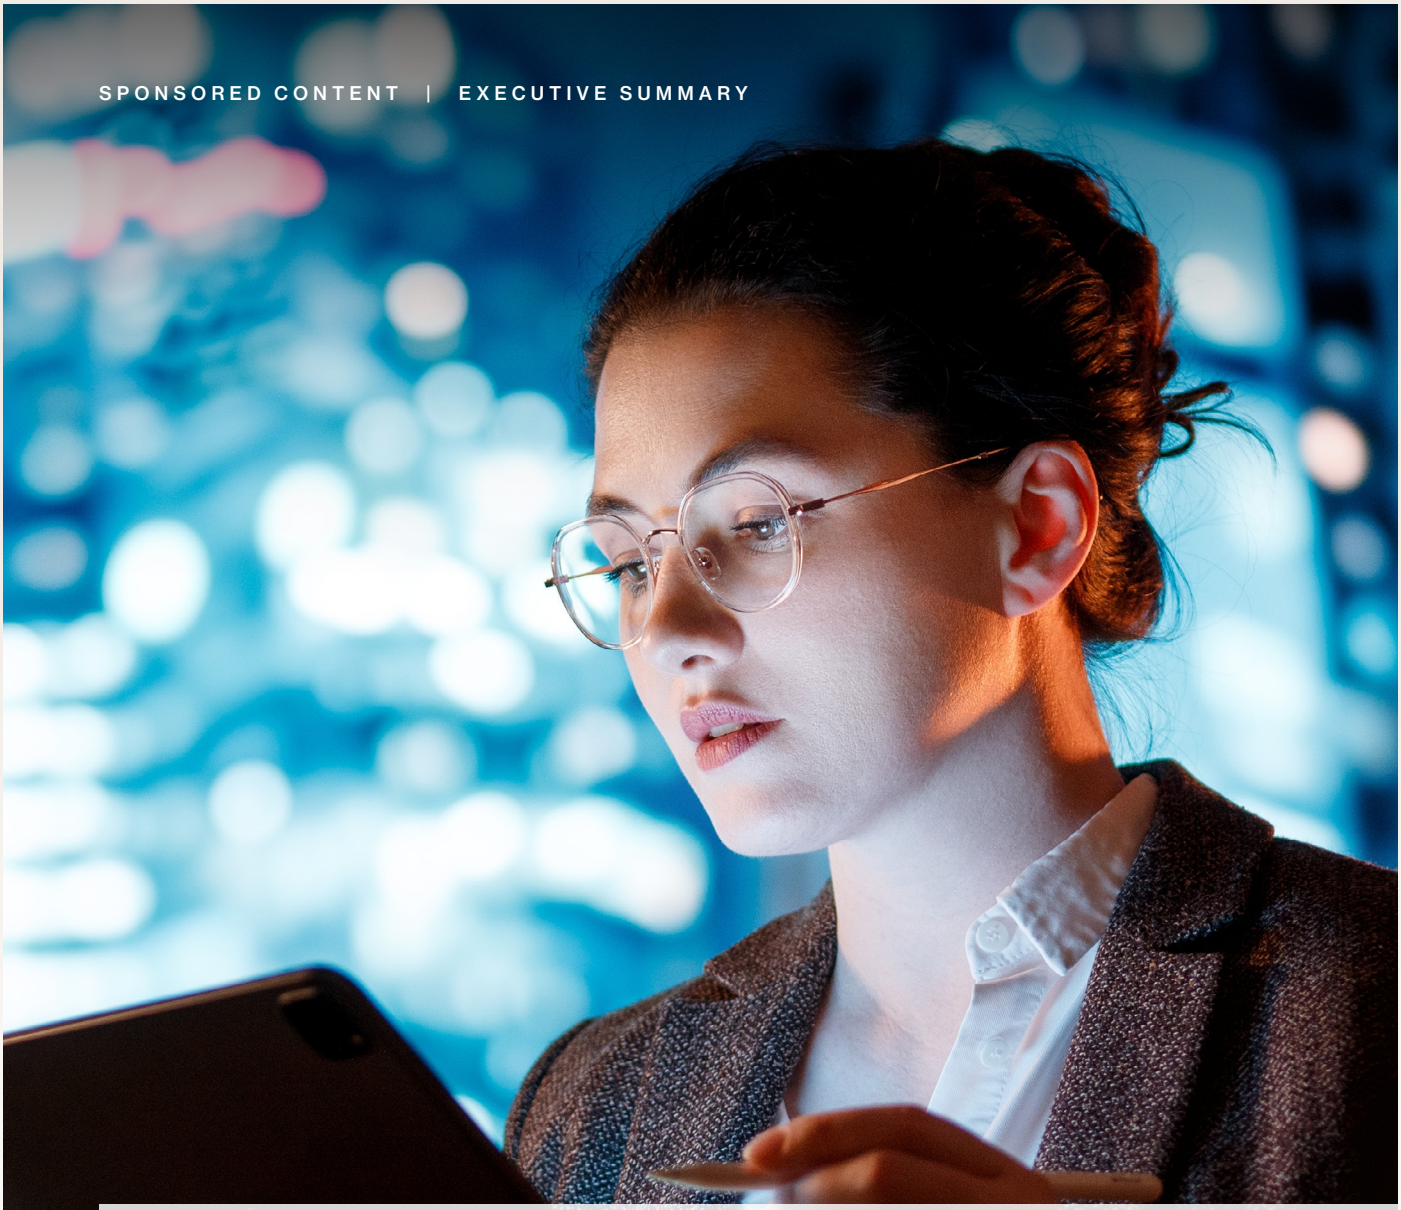

Executive summary:

The CIO's guide to unlocking scale with enterprise-grade GenAI

CIO

SPONSORED BY

NUTANIX

Market Pulse

Generative AI (GenAI) is no longer confined to experimentation. It is now a production technology that IT must operate reliably and securely at scale to deliver business value that directly impacts cost, risk, operational resilience, and long-term competitiveness.

For executive teams, the shift from pilots to production is material. GenAI is increasingly viewed as an enterprise-grade capability, on the same level as ERP, CRM, and other mission-critical systems. In a new Foundry study, 96% of the surveyed IT leaders acknowledged that GenAI is important to their organization's business strategy, and nearly as many see it as a competitive differentiator.

The strategic mandate is clear, yet the operational implications are still unfolding and the challenges are growing. To succeed at scale, GenAI initiatives must align with the same repeatable enterprise-wide practices that many organizations are using across their hybrid IT environments.

GenAI is now an enterprise priority

GenAI adoption is gaining traction across mission-critical production environments (see Figure 1). What began as proof-of-concept experimentation is rapidly moving into software-as-a-service (SaaS) applications, enterprise workflows, and customer-facing processes.

Two-thirds of the survey respondents reported that their organization is using GenAI in SaaS applications and embedding the technology into enterprise workflows. These deployments tap into a mix of structured, unstructured, and object data — virtually the entire data ecosystem that courses through an enterprise (see Figure 2).

About the research

Foundry conducted an online survey, sponsored by Nutanix, of 301 CIOs and senior IT leaders in Asia-Pacific, North America, and Western Europe. The respondents work at organizations with 1,000+ employees, and their enterprises have either piloted or implemented generative AI solutions. The survey was conducted between December 7 and December 23, 2025.

Market Pulse

Figure 1 | **GenAI gains traction across production environments**

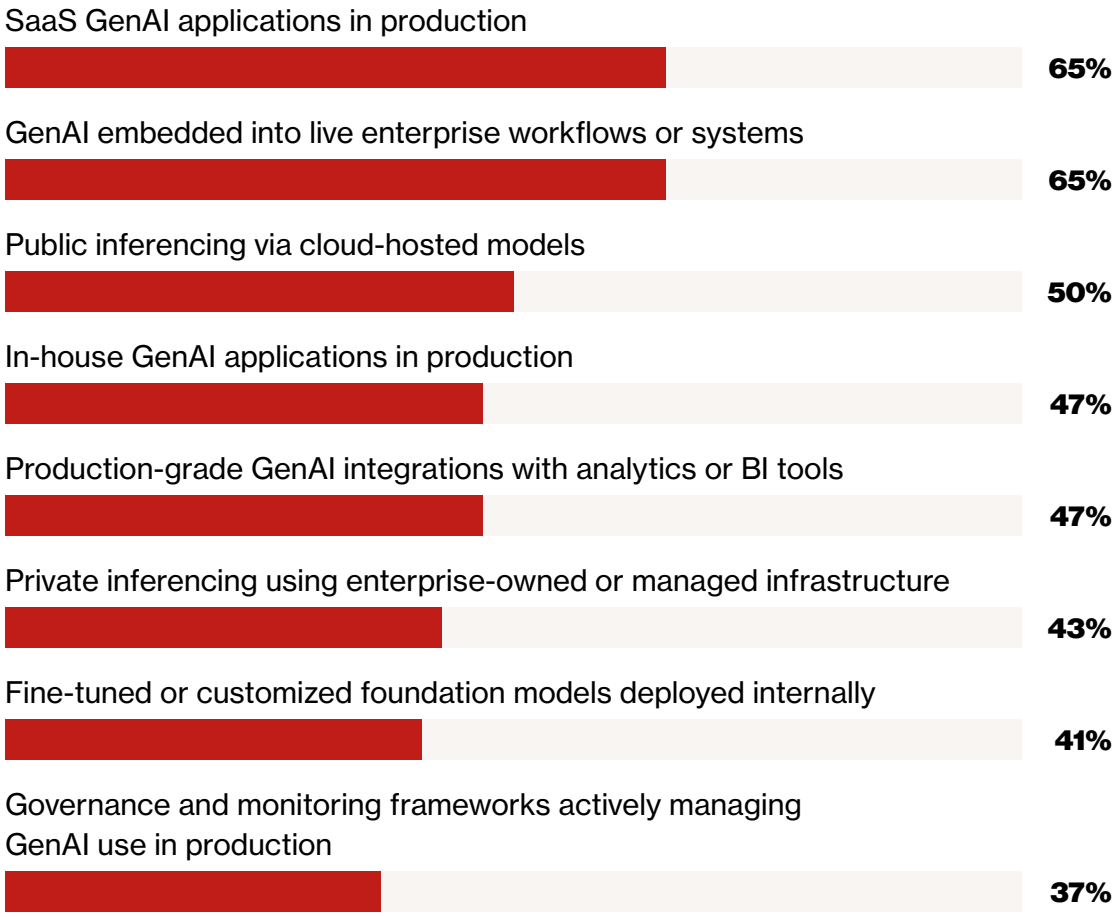

Figure 2 | **GenAI projects access multiple data types**

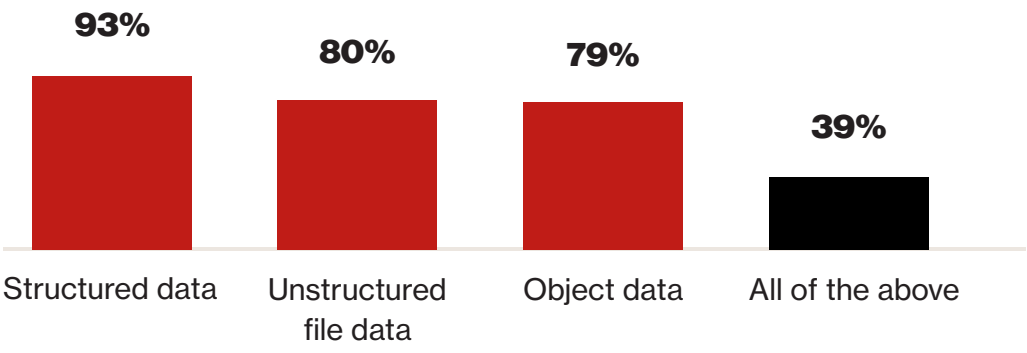

SOURCE: FOUNDRY

Market Pulse

GenAI's enterprise reach is broad and deepening. Most organizations are running or plan to run GenAI applications across hybrid cloud environments, including public clouds, private clouds, data centers, and edge locations. With enterprise data highly distributed across multiple locations, AI workloads increasingly span heterogeneous environments that must simultaneously support performance, reliability, and compliance.

The expected business payoff from GenAI investments reflects the C-suite's ambitions for the technology. The survey found that IT leaders anticipate outcomes such as:

- **Productivity and efficiency gains**
- **Improved customer experience**
- **Revenue growth**
- **Enterprise-wide transformation of business processes**
- **Innovation and new product development**

Enterprises are investing accordingly. Budgets are expected to increase to support these outcomes, and a strong majority of the leaders anticipate measurable ROI within the next 12 to 24 months.

Early production success, however, does not equate to enterprise operational readiness at scale. As GenAI expands across environments and workflows, it begins to reshape the IT operating model required to run, govern, and secure these systems consistently.

Successful leaders increasingly view GenAI as any other business-critical application — albeit an extremely powerful one — that requires enterprise-grade infrastructure, governance, and skills to support deployment and oversight across different environments.

C-suite implication

GenAI is already embedded in enterprise operations and is being funded with high expectations for ROI and competitive advantage. The next phase will test whether organizations can operate AI with the same rigor and resilience as core enterprise.

Infrastructure, risk, and the reality of scale

As GenAI adoption accelerates, structural challenges are becoming more visible. Many organizations are still working to scale

Market Pulse

Figure 3 | **The top infrastructure challenges with scaling GenAI**

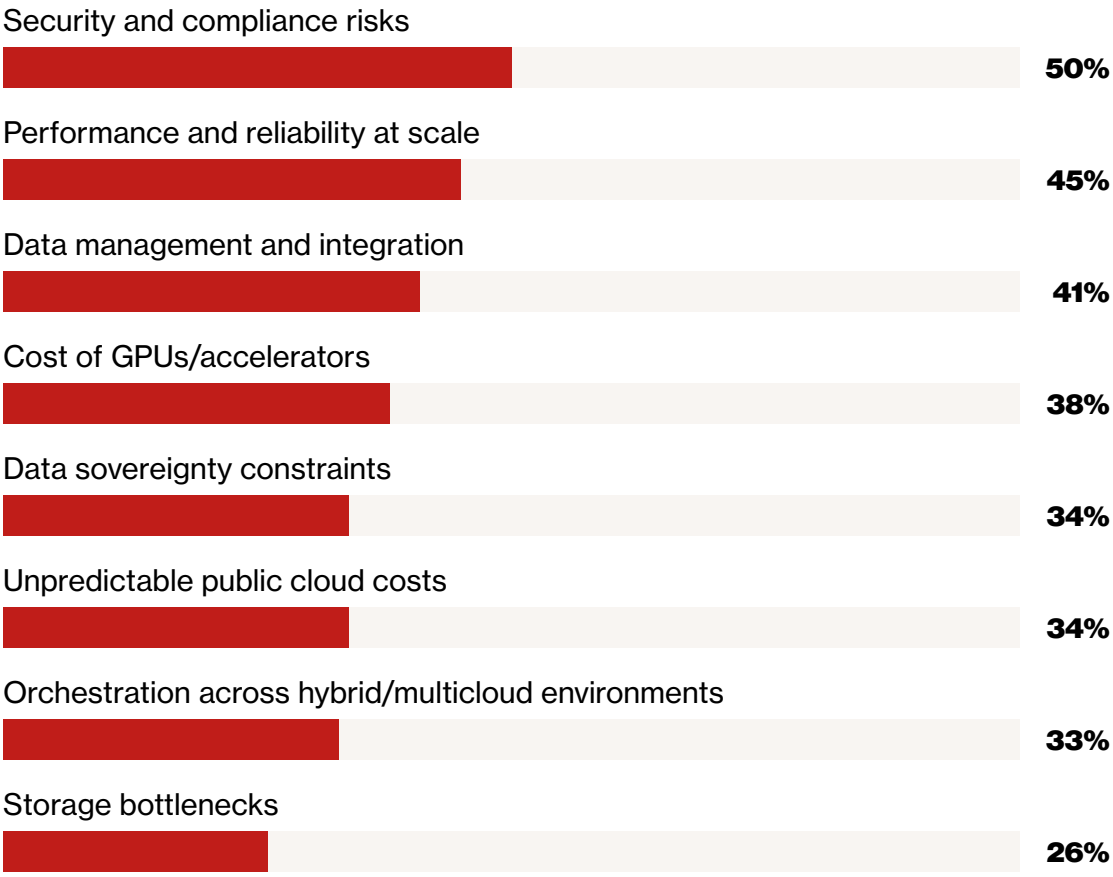

SOURCE: FOUNDRY

environments that can support AI workloads reliably, securely, and cost-effectively.

Asked about the top infrastructure challenges their organization faces in scaling AI, respondents cited security and compliance, performance and reliability, and data management and integration, among other factors (see Figure 3).

In parallel, leaders identified additional inhibitors to achieving GenAI goals, including:

- **Data risk:** Most respondents (94%) reported confidence in managing GenAI-related data privacy and compliance risks. Yet sensitive data used for AI often must remain within national or organizational sovereignty boundaries, and many platforms

Market Pulse

lack the controls to enforce residency and access policies consistently across environments. Organizations may also lack governance tools to secure AI pipelines and enforce policies across data, model, and compute layers. As GenAI scales, policy enforcement must become continuous, not episodic.

- **Skills readiness:** Demand for AI-critical skills – including AI governance and agentic workflow design – is outpacing supply. Talent shortages introduce scaling and resilience risks, particularly as AI systems evolve into more autonomous and integrated workflows. With a widening skills gap, enterprises must design AI infrastructure and governance models that are repeatable and automated, reducing dependency on scarce expertise.

- **Unpredictable costs:** IT leaders have long struggled to forecast and manage pay-as-you-go cloud costs. GenAI introduces even more variability, as advanced AI systems create overhead in tokens, latency, energy consumption, and compute intensity. Fragmented solutions can increase redundancy and waste as complexity rises and visibility decreases. Cost volatility, combined with unclear return on investment (ROI) from existing GenAI environments, increases executive scrutiny.

C-suite implication

Left unchecked, these issues translate directly into business risk – regulatory exposure, unstable service delivery, escalating operating costs, and erosion of executive confidence in AI-driven outcomes. Sustainable success requires acknowledging and addressing these tensions before they undermine business confidence.

From point solutions to platforms

As AI scales, fragmentation increases executive exposure across cost, governance, and operational resilience. The research signals a fundamental shift: Enterprises must move from isolated point solutions to a unified operating model.

This platform approach is not about consolidating tools. It is about consistency and simplifying AI operations so existing IT teams can deploy, govern, and operate GenAI by using consistent workflows and familiar operational models.

Market Pulse

Treating AI like any other enterprise application will help ensure that AI systems are deployed strategically and run with resilience, Day 2 operations, and security. Without this shift, GenAI initiatives risk fragmenting further into siloed environments that increase cost volatility, weaken accountability, and limit the ability to deliver predictable business outcomes.

Three core components are critical for an effective enterprise AI operating model:

- **Support across hybrid environments:** Enterprise data is distributed across on-premises data centers, public clouds, and the edge. AI workloads must run where data, latency, and sovereignty requirements dictate. Hybrid-ready architecture is no longer optional — it is foundational to performance and resilience.

- **Consistent governance and built-in risk management:** Because sensitive data used for AI modeling often must remain within specific boundaries, AI platforms must enforce data residency, respect regional sovereignty, and enable control over where and how data is processed. Governance policies must address data lineage, access controls, usage monitoring, and life cycle policy enforcement across environments.

- **Predictable costs and demonstrable ROI:** Fragmented point solutions limit visibility and increase volatility. A unified platform with end-to-end observability simplifies performance and cost optimization, enabling leaders to identify bottlenecks, align workloads, and control expenses as AI expands. Automated Day 2 operations improve efficiency and reduce manual intervention, strengthening the likelihood of sustained ROI.

C-suite implication

AI scalability depends less on model selection and more on operational architecture. Enterprises that establish unified platforms, embedded governance, hybrid flexibility, and cost visibility will convert GenAI investments into repeatable, resilient value.

Market Pulse

Becoming a TCO-focused innovation enabler in the GenAI era

GenAI adoption is accelerating into mission-critical production environments. It is shifting from experimentation to an enterprise operating strategy that directly impacts cost structure, resilience, governance, and long-term competitiveness.

Executive teams must bring the same discipline to GenAI environments that they apply to core systems, but at greater scale and velocity.

The Foundry report "The CIO's guide to unlocking scale with enterprise-grade GenAI" explores additional dimensions shaping enterprise AI strategy, including:

- **Investment acceleration and ROI expectations**
- **Ownership and accountability models for enterprise GenAI**
- **The rise of agentic AI and its governance implications**
- **Why GenAI collapses the traditional Day 0/Day 2 operational divide**

Explore the findings and analysis, and define your next step toward enterprise-scale GenAI.

Download the full report: "The CIO's guide to unlocking scale with enterprise-grade GenAI."

CIO

SPONSORED BY

NUTANIX