

SPONSORED CONTENT | WHITE PAPER

Market  
Pulse

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

Recent research offers insights into how IT leaders are preparing for the next phase of AI.



CIO

SPONSORED BY

NUTANIX

Generative AI (GenAI) has moved from experimentation into production. IT must now operate the technology securely and reliably to support enterprise-wide execution and measurable outcomes. This shift elevates it to a strategic operating and architectural priority that shapes cost predictability, governance, business continuity, regulatory exposure, and long-term competitiveness.

CIOs and enterprise leaders are beginning to view GenAI as an enterprise-grade capability on the same level as ERP, CRM, and other mission-critical systems, according to a new Foundry Market Pulse study. Virtually all senior IT leaders in the survey (96%) acknowledge that GenAI is important to their organization's business strategy. Almost as many agree that it is a competitive differentiator (see Figure 1).


This strategic mindset makes sense, given CIOs' mandate to ensure that their burgeoning GenAI investments scale, deliver timely payback to the business, and provide a foundation for resilience and sustainable success. As a result, IT leaders should operationalize GenAI as a standardized platform designed to scale across hybrid environments and governed by the same principles that drive resilience, predictability, and accountability.

Forward-thinking and successful leaders view the technology as being on par with any other business-critical application


**Figure 1**  
**GenAI is now an enterprise priority**




**GenAI is important to our business strategy in the next 12 months**



**Our business views GenAI as a competitive differentiator**



**GenAI is on par with ERP and CRM**



SOURCE: FOUNDRY

(albeit an extremely powerful one) that requires enterprise-grade infrastructure, governance, and skills to support deployment and oversight across different environments. It must be engineered, governed, and operated as enterprise infrastructure.

## GenAI is now an enterprise priority

GenAI adoption is gaining traction across mission-critical production environments. For example, two-thirds of the survey respondents said their organization is using GenAI in software-as-a-service (SaaS) applications and has embedded it 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).

Another indicator of the broad and deepening reach of GenAI workflows: Most organizations are running or plan to run GenAI applications across hybrid cloud environments, including public clouds, private clouds, data centers, and edge locations. As CIOs are quickly learning, this type of scale raises myriad integration, security, sovereignty, and performance challenges.

For IT leaders who can overcome these issues, the payoff will be worth it. IT leaders anticipate a range of positive outcomes from their GenAI investments, including:

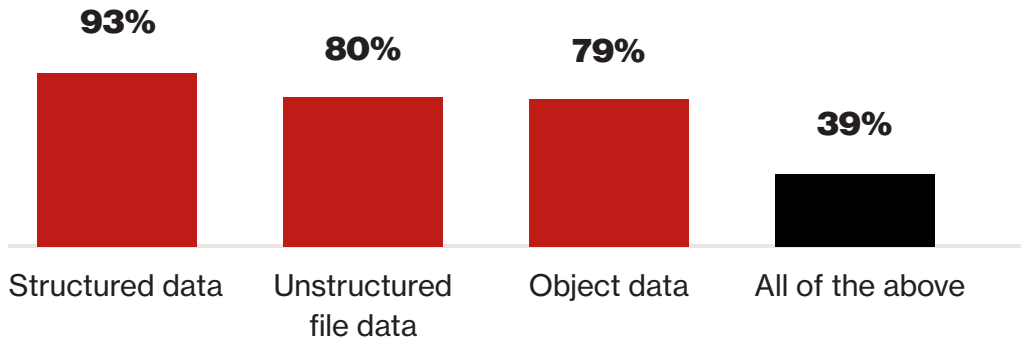
- Productivity and efficiency gains
- Improved customer experience
- Revenue growth
- Enterprise-wide transformation of business processes
- Innovation and new product development

Early production success, however, does not necessarily indicate enterprise readiness or long-term sustainability at scale.

“There’s definitely a positive outlook about how GenAI can impact strategic areas in enterprises,” says industry veteran Harsha Kotikela, senior director of product and solutions marketing at Nutanix. “But just like with any new tech, many organizations are still in the early stages of their journey. The conversation among CIOs is now shifting from excitement to ‘How do we take what we’ve learned from all the proofs of concept and build the foundational pieces that will allow us to scale?’”

**Figure 2 | GenAI has broad reach across business and IT environments**

**Types of data accessed by GenAI projects**



**GenAI adoption in production environments**


SaaS GenAI applications in production




GenAI embedded into live enterprise workflows or systems




Public inferencing via cloud-hosted models




In-house GenAI applications in production




Production-grade GenAI integrations with analytics or BI tools




Private inferencing using enterprise-owned or managed infrastructure



Fine-tuned or customized foundation models deployed internally



Governance and monitoring frameworks actively managing GenAI use in production



SOURCE: FOUNDRY

## Challenges: Risk, infrastructure, and scale

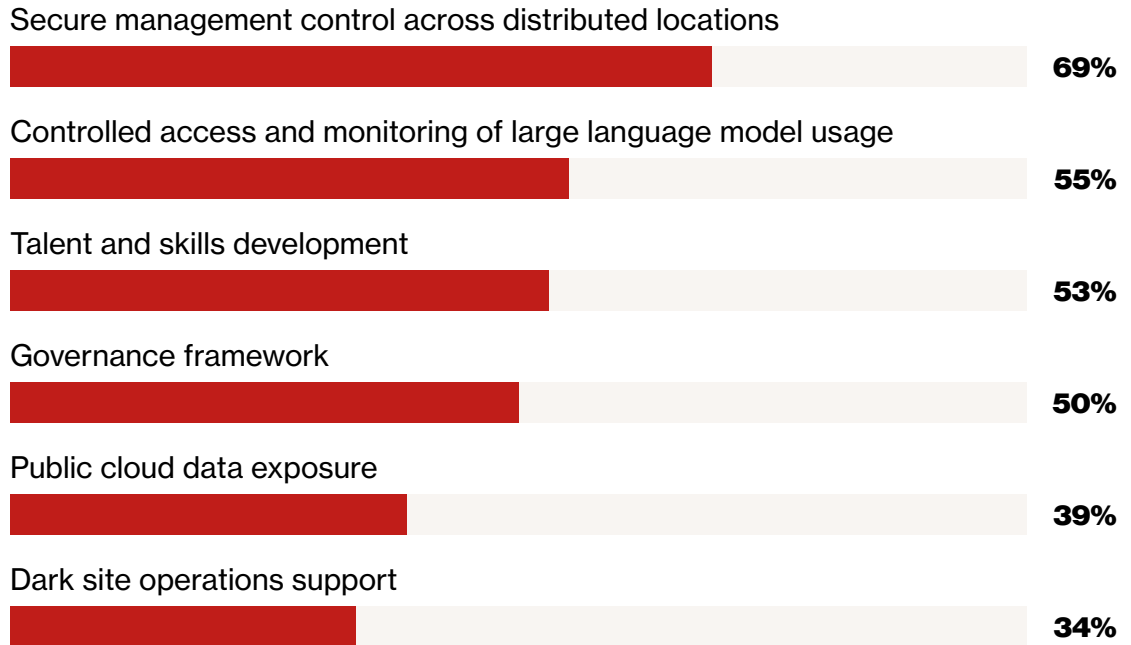
Asked about the infrastructure challenges their organizations face in scaling AI, respondents cited security and compliance, performance and reliability, and data management integration as the top three.

In addition, IT leaders called out significant challenges to achieving their GenAI goals, including:

- **Data risk:** Most respondents (94%) said they're at least moderately confident in their ability to manage GenAI-related data privacy and compliance risks. They may be underestimating the effort needed to upgrade platforms, policies, and tools in support of AI's data governance requirements. For example, sensitive data used for AI often must remain within national or organizational boundaries. GenAI increasingly requires sovereign execution – ensuring that data, models, and processing remain within defined jurisdictional or corporate controls. Yet many platforms lack the transparency and enforcement capabilities to meet these requirements. Organizations may also lack integrated governance frameworks to secure AI pipelines and enforce policies across data, model, and compute layers.
- **Skills readiness:** Demand for AI-critical skills, including AI governance and agentic workflow design, is outpacing supply, leaving enterprise leaders unprepared to address the integration and orchestration requirements of agent-based systems. Talent shortages also introduce scaling and resilience risks. With a widening skills gap, CIOs will need to build AI infrastructure and governance that is repeatable and automated, thus reducing their dependency on scarce specialized AI talent.
- **Unpredictable costs:** IT leaders have long struggled to improve how they forecast and manage the costs of pay-as-you-go cloud services. AI introduces even more unpredictability into IT cost management. Advanced AI systems can create significant overhead in tokens (units of data processed by AI models during training and inference), latency, and energy. Fragmented solutions can create redundancy and waste as complexity rises and visibility decreases. Unpredictable costs will drive CIOs to improve visibility, optimization, and cost-tracking capabilities for AI tools and platforms across hybrid environments.

**Figure 3**

**IT leaders seek to reduce critical GenAI privacy risks**



\*Which areas do you consider most critical to reduce GenAI privacy risk? (Rank Top 3)

SOURCE: FOUNDRY

Structural challenges will only increase as GenAI expands across multiple teams, systems, and environments.

Left unchecked, this complexity will translate explicitly into business risks, including regulatory exposure; unpredictable operating costs; service instability; and, at a more existential level, the erosion of executive confidence in AI-driven outcomes.

“The pace of change is so fast with AI that if you’re not

keeping up, you will essentially lose your competitive edge,” says Kotikela. “Data and infrastructure are strategic to the business, and if you get those wrong with AI, the impact will be amplified.”

IT leaders will need to balance early optimism with production success and scalable enterprise readiness. Achieving GenAI scale – and capturing the expected benefits – will require a renewed focus on the foundational components that are critical to long-term success.

## From point solutions to platforms

Fragmented AI solutions increase executive exposure across cost, governance, and operational resilience. To reduce risk and provide a reliable path to success, CIOs should approach AI like any other enterprise application, ensuring that AI systems are deployed strategically and run with resilience, Day 2 operations, and security.

The best way to reduce risk and increase resilience is to establish a single enterprise operating model for AI that delivers consistent governance, cost predictability, and resilience. This platform approach is not about consolidating tools; it's about simplifying AI operations so existing IT teams can deploy, govern, and operate GenAI with consistent workflows and familiar operational models.

Without a platform, GenAI initiatives risk fragmenting further into siloed point solutions that increase cost volatility, weaken accountability, and limit the ability to deliver predictable business outcomes.

Among the key considerations:

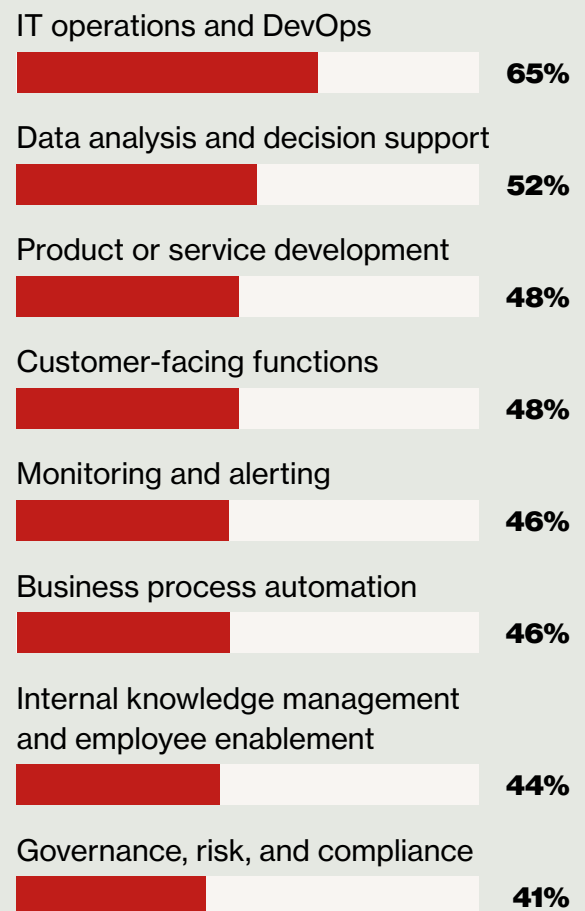
- **Support across hybrid environments:** AI workloads are pushing infrastructure to its

## Agentic AI raises the stakes for enterprise readiness

GenAI agents are moving beyond customer-facing chatbots into core operations. Agentic AI represents the next phase of enterprise AI adoption – and a step-change in complexity. Unlike single-purpose models, agentic systems involve multiple AI components that interact autonomously, passing data, triggering actions, and making decisions with limited human intervention.

**Figure 4**

### CIOs plan to deploy GenAI agents across the business over the next 12 months



SOURCE: FOUNDRY

limits. Performance, reliability, and resilience are critical to turning AI into a durable source of business value. However, because the enterprise data that fuels AI is highly distributed across multiple locations – including on-premises data centers, public clouds, and the edge – CIOs will need to deploy a hybrid architecture that enables AI workloads to run wherever data, latency, and sovereignty requirements dictate.

- **Consistent governance, with built-in risk management and improved visibility, accountability, and policy enforcement across environments:** Because sensitive data used for AI modeling often must remain within national or organizational boundaries, AI platforms must enforce data residency, respect regional sovereignty, and enable control over where and how data is processed. Deploying large language models (LLMs) in an enterprise environment requires robust validation pipelines, ongoing performance monitoring, and alignment with risk and compliance policies. AI governance policies will need to address data lineage, access controls, usage monitoring, and life cycle policy enforcement.

← Continued from sidebar on page 7

This escalation of agentic AI deployments fundamentally alters enterprise risk. Each agent introduces new governance requirements: which data it can access, which systems it can influence, and how its actions are monitored and constrained. As agents begin working together in chained workflows, risk compounds across data exposure, cost predictability, performance, and accountability.

Many enterprises are eager to explore enterprise AI initiatives to accelerate automation and decision-making. However, deploying agents without a mature operating model often leads to fragmented controls and blind spots. Traditional guardrails built for static applications are insufficient when AI systems dynamically evolve and interact.

For CIOs, enterprise AI reinforces a critical lesson: Future innovation depends on today's foundations. Enterprises that invest early in governance, observability, and standardized AI operations will be positioned to adopt agents safely and confidently. Those that do not may find that innovation outpaces control, introducing risks that scale faster than value.

■ **Predictable costs and the ability to demonstrate ROI from GenAI investments:**

Fragmented point solutions increase cost volatility and limit the ability to deliver positive business outcomes. By comparison, a platform with end-to-end observability simplifies both performance and cost optimization, making it easier to identify bottlenecks, align workloads, and control costs as AI grows. Automated Day 2 operations for GenAI applications across hybrid environments improve efficiency and help deliver ROI.

“An AI operating environment integrates the full stack needed to operationalize enterprise AI – from the physical infrastructure layer to orchestration, platform services, and centralized management,” says Kotikela. “The goal is to build a platform that delivers best-of-breed capabilities at every level of the stack.”

**Become an innovation enabler in the GenAI era**

As GenAI adoption moves rapidly into a variety of mission-critical production environments, including

SaaS applications and enterprise workflows, it is shifting from an innovation initiative to an enterprise operating strategy that directly impacts cost structure, resilience, governance, and long-term competitiveness. That’s why CIOs must bring the same discipline to GenAI environments that they apply to core systems – but with greater scale and velocity.

Defining the right use cases for GenAI is important for framing an operating model designed to automate and optimize processes and workflows. The operating model serves as a blueprint for updating data pipelines and infrastructure to support AI workloads and applications.

A platform approach helps CIOs move GenAI from pilot to production across the enterprise, providing a turnkey, resilient, secure, scalable environment that accelerates innovation and delivers positive business outcomes. That’s a win for any CIO.

# GenAI investments are growing, but what about payback?

Core IT spending on GenAI is expected to grow dramatically, from \$40.1 billion in 2024 to \$151.1 billion in 2027, according to IDC,<sup>1</sup> which calls this period a critical AI buildout phase for enterprises worldwide.

The Foundry survey found that most IT leaders will increase their GenAI budget over the next two years, by an average of 12% over that period.

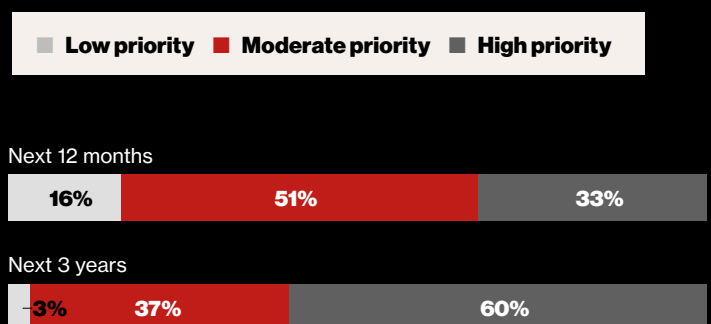
Also, enterprise leaders anticipate fairly quick payback from these investments. More than eight in 10 (84%) are at least moderately confident that they will achieve a positive return on investment (ROI) in the next 12 months. And virtually all of them (97%) are confident that their GenAI initiatives will deliver measurable business value over the next two years (see Figure 5).

Asked which capabilities will most enable successful at-scale GenAI adoption over the next three years, CIOs and enterprise leaders cited:

1. Ability to demonstrate predictable ROI from GenAI investments
2. Speed of innovation and experimentation
3. Access to skilled talent
4. Infrastructure resilience and performance
5. Governance and responsible GenAI practices

As ROI pressure increases, CIOs will likely need to shift toward a platform-centric approach versus point solutions for their GenAI workflows and applications. Scalable platforms enable IT leaders to deliver autonomous workflows, data sovereignty, and end-to-end observability. These key contributors to GenAI deployments will increase the likelihood of timely, sustainable payback.

**Figure 5**  
**IT leaders have high expectations for GenAI ROI**



SOURCE: FOUNDRY

<sup>1</sup> IDC, "Generate Growth in Your Markets with the GenAI Opportunity."

## Who owns enterprise GenAI?

## Why ambiguity is a hidden risk

As GenAI scales across the enterprise, one issue quietly surfaces inside many enterprises: unclear ownership. Survey data shows that responsibility for GenAI governance is often distributed across IT, security, and AI or data executives – by either design or default. This shared model reflects the cross-functional nature of AI, but it also introduces risk.

Unlike traditional applications, GenAI systems span data pipelines, models, infrastructure, and automated decision-making. Without clear accountability, enterprises struggle to enforce consistent standards for security, compliance, cost controls, and performance. Gaps emerge, not because teams lack intent but because no single executive owns the end-to-end operating model.

For CIOs, this ambiguity presents both a challenge and an opportunity. As GenAI becomes a core part of enterprise infrastructure, leadership expectations shift toward centralized governance paired with decentralized execution. IT leaders don't need to control every AI initiative, but they are increasingly expected to define the operating framework: how AI is deployed, governed, monitored, and scaled safely across the organization.

Enterprises that resolve GenAI ownership early gain more control – as well as speed, confidence, and credibility – as AI expands into mission-critical workflows.

# A strategic framework for operationalizing GenAI

For decades, enterprise IT has distinguished between Day 0 deployment and Day 2 operations. GenAI collapses that distinction. Although many organizations focus on launching models and pilots, long-term success depends on how AI systems behave after they go live across environments.

GenAI introduces new operational demands that exceed traditional application management. CIOs must account for continuous model monitoring, data lineage tracking, agent behavior, and token-based cost variability. Observability must extend beyond uptime to include performance, accuracy, usage patterns, and policy compliance. Security expands from perimeter defense to model access controls and machine-to-machine trust.

## What CIOs mean by enterprise-ready AI

- **Predictable economics**
- **Governed data flows**
- **Hybrid execution**
- **Day 2 automation**
- **Business-aligned outcomes**

Day 2 operations for GenAI also require automation. Manual interventions that worked for earlier workloads quickly become unsustainable as AI systems multiply and evolve. Enterprises need built-in resilience, automated remediation, and consistent governance across hybrid environments.

The lesson for CIOs is clear: GenAI cannot be treated as a special-case workload. It must be operated with the same rigor as ERP or CRM, augmented by new controls suited to AI's scale, speed, and autonomy. The enterprises that master Day 2 operations will be the ones that turn AI ambition into durable advantage.

**Visit [Nutanix Executive Focus](#) for more executive insights and thought leadership on this subject.**

## About the research

Foundry conducted an online survey, sponsored by Nutanix, of 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.

## About Nutanix

Nutanix is a hybrid multicloud computing leader, offering organizations a unified software platform for running applications, deploying enterprise AI workloads and managing data anywhere. With Nutanix, organizations can simplify operations for traditional and modern applications, freeing them to focus on business goals. Trusted by more than 30,000 customers worldwide, Nutanix helps empower organizations to transform digitally and power hybrid multicloud environments consistently, simply, and cost-effectively. Learn more at [www.nutanix.com](http://www.nutanix.com).

CIO

SPONSORED BY

NUTANIX