



How CIOs turn AI pressure into enterprise AI that scales

A roadmap for making AI predictable,
governable, and usable at scale

The challenge is no longer adopting AI. It's making AI work reliably, predictably, and securely across the enterprise.

AI has elevated the CIO's role to a new level of strategic visibility. Technology leaders are now expected to explain how AI will spur competitive advantage, operational resilience, and long-term enterprise performance. In many organizations, AI has moved from a technical curiosity to a priority board-level agenda item almost overnight.

Yet current technology environments were rarely designed to provide the support needed to enable AI to deliver these business outcomes.

Most enterprises still operate technology ecosystems shaped over decades. Core transactional systems contain deeply embedded business logic. Enterprise platforms carry years of customization. Integration architectures evolved organically across business units. Data environments support reporting and analytics rather than continuous orchestration across multiple systems.

This mismatch between ambition and operating reality explains why many organizations experience the same pattern. Early AI pilots demonstrate promising results. Leadership asks for the capability to scale. As initiatives move beyond controlled environments, deeper constraints appear, integration complexity increases, governance questions multiply, data readiness issues surface, and infrastructure costs become difficult to predict.

At that point, the conversation shifts. The question is no longer, "Where can AI help?"

The real question becomes, "What has to change so AI can operate reliably across the enterprise?"

Organizations that successfully answer that question do not simply run more experiments. They strengthen the structural conditions that allow AI to become a repeatable enterprise capability.

Diagnostic: Are the structural conditions for enterprise AI actually in place?

Many organizations measure AI progress by counting pilots or deployed tools. A more revealing assessment examines the environment surrounding those initiatives. Technology leaders often start with a few practical questions:

- Can the organization clearly explain where AI costs sit and where AI value appears across the business?
- Do successful AI pilots have a defined path to enterprise deployment?
- Are AI tools changing workflows and decision processes, or simply being deployed?
- Could the enterprise explain its AI governance model clearly to regulators tomorrow?
- If infrastructure conditions shifted suddenly, could the architecture adapt without major redesign?

If several of these questions are difficult to answer confidently, then it means the barrier to AI progress is rarely failed experimentation. It is enterprise readiness.

Why scaling AI exposes deeper structural problems

AI introduces a different operating model than most enterprise environments were built to support.

Models interact with multiple systems simultaneously, depend on reliable data across distributed repositories, and must operate within governance frameworks capable of tracking algorithmic decisions. As a result, organizations encounter what CIOs increasingly recognize as the pilot-to-production gap.

Here's a typical scenario: An AI pilot demonstrates meaningful AI value, leadership expands the project, but then progress slows dramatically. This delay is because scaling reveals the hidden fractures in the enterprise foundation:

- **Data-readiness challenges:** Pilots often run on carefully curated datasets. Enterprise-scale AI requires access to a federated data landscape plagued by silos, inconsistent quality, and a lack of semantic definition, forcing data science teams into a perpetual cycle of data cleansing.

- **Integration complexity:** A successful pilot might connect to two or three systems. A production solution may need to interact with dozens of legacy and modern applications, each with its own application programming interface, security model, and performance quirks.
- **Governance requirements:** A controlled pilot operates with minimal oversight. A production system must navigate a complex web of regulatory compliance (like GDPR or CCPA), internal model risk management, bias detection, and explainability requirements that were never designed into the initial architecture.

In the final analysis, the pilot didn't fail. The enterprise environment simply wasn't designed to support it. To move beyond incremental gains, CIOs must build a comprehensive technology operating model that makes AI coherent, controlled, and repeatable.

The roadmap for enterprise AI that scales

Organizations that move beyond isolated AI experimentation typically focus on several structural priorities simultaneously. These priorities reinforce one another and gradually create an environment where AI capabilities can operate predictably across the enterprise. Strategy affects what gets funded. Funding logic shapes what can be modernized. Architecture readiness influences whether AI can scale without bespoke effort. Adoption determines whether any of that work produces visible value. Operational resilience determines whether the whole effort survives changing conditions.

In practice, leading CIOs tend to focus on the following five priorities.



Align technology strategy with enterprise priorities

When it comes to enterprise-wide AI deployment, CIOs need a defensible plan that reconciles business objectives, AI

ambitions, cost constraints, and operational risk into a phased execution roadmap. This plan starts by creating a foundational operating model to guide all AI initiatives.

The demand for AI often emerges simultaneously across multiple parts of the business. Product teams seek intelligent customer experiences. Operations leaders pursue automation. Finance groups explore predictive forecasting. Risk teams want improved monitoring and compliance insight.

Without a coherent strategy, CIOs often respond to these requests individually. Although this enables innovation to move quickly, it gradually fragments the technology environment. Each initiative introduces additional tools, integration pathways, and governance requirements that make the environment harder to manage.

Organizations that scale AI effectively begin by defining how AI capabilities will operate across the enterprise. Integration standards, identity frameworks, data governance principles, and platform priorities are clarified early so innovation occurs within consistent guardrails.

This strategic clarity allows technology investments to be sequenced deliberately and grounded in business priorities. Some initiatives move forward immediately, while others wait until architectural conditions make them easier to scale.

By defining the future-state operating model first—clarifying IT services, establishing a central AI Center of Excellence for governance, and standardizing technology patterns—the organization's design is built to support the strategy.

Outcome — CIOs achieve faster executive alignment through a defensible strategy that integrates spend control and risk reduction. This leads to measurable stability and cost transparency that unlocks subsequent modernization.

2

Build a credible narrative for AI investment

Another obstacle to enterprise AI progress is the difficulty of explaining its financial impact. Many

organizations still try to explain AI through a narrow ROI frame centered on immediate cost reduction or headcount savings. That framing is often too limited for the way AI creates value in practice. Instead, CIOs need to make technology spend explainable, controllable, and optimizable to create a self-funding engine for innovation.

AI value frequently appears incrementally, manifesting in faster analysis, improved decision quality, and shorter planning cycles. These gains are meaningful, but they rarely translate immediately into traditional financial metrics.

Meanwhile, the infrastructure required to support AI—cloud platforms, data engineering environments, integration services—often appears clearly in the IT budget. The result is a familiar tension: Technology costs are visible, but the benefits are harder to discern because they are distributed across the enterprise.

Leading CIOs address this challenge by broadening the investment narrative. Rather than focusing exclusively on short-term savings, they connect AI initiatives to enterprise outcomes such as decision speed, customer responsiveness, risk mitigation, and operational resilience.

Improving transparency around technology spending is equally important. When CIOs can explain what drives infrastructure costs and how those investments support enterprise capabilities, finance leaders gain greater confidence in long-term modernization efforts.

To make technology spend explainable and controllable, they master tech-asset management, connecting financial, contractual, and inventory data to optimize licenses and establish a robust FinOps practice for managing the costs of AI. This creates a “self-funding” engine for innovation.

Outcome — Spend becomes forecastable, shifting CFO conversations from reactive inquiries to planned, value-driven investments. Modernization and AI initiatives become fundable without perpetually needing new budget asks, and vendor risk visibility improves dramatically.

Three questions CIOs ask before major AI budget discussions

Technology leaders who successfully secure AI investment often prepare by answering three questions clearly:

- **Where has AI changed actual workflows or decision processes?**
- **What are the primary drivers behind infrastructure spending?**
- **Which earlier investments made current AI capabilities possible?**

Clear answers often transform funding conversations from defensive explanations into strategic discussions.



3

Strengthen architecture and data foundations

Infrastructure fragmentation remains one of the most common barriers to enterprise AI scale. As AI shifts

from isolated use cases to strategic applications, fragmented systems and brittle integrations create significant risk and complexity. An AI-ready architecture—modular, resilient, and data-centric—is required for repeatable value. Arriving at this state means moving beyond project-based data prep to building a scalable data fabric or mesh with clear data ownership and trusted data products.

This legacy of organically evolved integration makes each new AI initiative a heavy engineering lift. Within these environments, each AI initiative requires significant engineering work to connect systems, validate security assumptions, and manage operational dependencies.

Organizations that scale AI successfully address this challenge by establishing shared architectural patterns. Integration standards become consistent across platforms, identity and access frameworks align across systems, and data governance models clarify ownership and lineage.

Over time, these improvements transform the environment. Instead of rebuilding infrastructure for every initiative, organizations create reusable foundations that support continuous innovation.

Outcome — AI becomes operational, moving beyond fragile pilots to repeatable, production-grade patterns with defined guardrails. Delivery speed increases without increasing fragility, and costs become governable even as AI scales exponentially.

4

Turn AI adoption into an operating model

Even with the right technology and architecture, human adoption and enablement are crucial for realizing value.

Cultural resistance, persistent talent gaps, and a lack of executive fluency can hinder any transformative effort. Even when technology environments improve, many organizations discover that adoption remains uneven.

AI tools may be widely available, yet real workflow changes occur slowly because technology alone doesn't reshape operating behavior. Teams continue using familiar processes, managers hesitate to redesign workflows around new capabilities, and technology becomes a supplemental tool rather than a transformative one.

In overcoming this issue, generic, one-size-fits-all training is ineffective. Successful CIOs treat adoption as an operating model challenge rather than a rollout exercise, implementing targeted enablement programs—like boot camps and embedded expert support—that are integrated directly into real-world workflows.

Lasting transformation is business-led, so responsibility for AI-driven process redesign must be delegated to functional leaders, with the IT function acting as an enabler providing platforms and guardrails.

This fosters a culture of AI empowerment. Lasting transformation is business-led, so responsibility for AI-driven process redesign must be delegated to functional leaders, with the IT function acting as an enabler providing platforms and guardrails. Successfully achieving this outcome requires developing executive AI fluency through dedicated education so the C-suite can become active, informed sponsors.

Executive education programs help leaders understand how AI changes decision-making. Technology expertise is embedded within business units so experimentation occurs close to operational needs. Governance frameworks allow innovation while maintaining consistent standards.

In this setting, adoption is measured through workflow improvement rather than deployment counts. When AI influences how work is performed, its value becomes visible much more quickly.

Outcome — CIO capacity returns, allowing leadership to shift from firefighting to executing strategic priorities. Operations stabilize, leading to fewer surprises, and the organization develops a resilient, AI-ready workforce capable of leading continuous process redesign.

5

Strengthen operational resilience

CIOs today find themselves grappling with internal challenges while navigating a rapidly evolving

external environment where economic, regulatory, and geopolitical forces reshape priorities.

Economic uncertainty shortens planning horizons, forcing CIOs to prioritize initiatives that demonstrate measurable progress within the current year rather than focusing solely on long-term potential. AI regulatory frameworks evolve across jurisdictions, elevating “responsible AI” from a policy discussion to an operational design imperative. Geopolitical dynamics influence vendor ecosystems and data sovereignty decisions, placing increased pressure on CIOs to demonstrate control over data location, vendor dependencies, and security partner geography.

Even as these external forces exert themselves on the organization, teams are often consumed by addressing internal day-to-day challenges and maintaining operations.


Organizations that adapt successfully invest in resilience before disruption occurs. They improve visibility into system dependencies, which is crucial because CIOs cannot effectively reduce operational risk when dependencies are unknown or architectures are brittle. They also strengthen testing and deployment practices, reducing risk by ensuring improved test coverage and standardized run operations. They clarify governance policies around vendor relationships and data residency, establishing operationalized governance that includes strict permissions, controls, and monitoring for AI-enabled workflows.

These efforts rarely attract the same attention as launching new AI initiatives, yet they significantly increase the organization’s ability to sustain innovation. By reframing foundational work as a business necessity—linking it to risk reduction, cost predictability, and speed—CIOs can build a compelling case for investment that resonates with C-suite executives and the board, enabling sustainable innovation.

Outcome — Operational stability, combined with a modern architecture, improves visibility into system dependencies and increases the organization’s ability to sustain innovation through disruption.



Operationalizing the roadmap: the capability architecture behind scalable AI



The roadmap outlined above is conceptually simple but operationally demanding. Each priority requires coordinated changes across strategy, architecture, governance, cost management, and engineering. Organizations that make the fastest progress rarely address these dimensions sequentially. Instead, they build an integrated capability architecture that allows them to move forward on multiple fronts simultaneously.

This architecture typically combines several complementary capabilities:



Technology strategy and operating model design helps CIOs define how AI initiatives align with enterprise priorities and how technology investments should be sequenced. Establishing a clear operating model also clarifies decision rights, governance structures, and service models that determine how innovation progresses across the enterprise.




Technology asset management and cost transparency provide the financial clarity required to sustain modernization efforts. By connecting financial, contractual, and inventory views of technology assets, organizations gain visibility into infrastructure spending and can identify opportunities to optimize licensing, reduce redundancy, and reinvest savings into higher-value initiatives.



Architecture and engineering modernization strengthens the technical foundations that allow AI to scale. Standardized integration patterns, clearer data governance frameworks, and improved observability practices reduce the engineering overhead required to launch new initiatives and make enterprise environments easier to manage.



Lifecycle modernization platforms such as KPMG Ignite | Blaze address one of the most difficult challenges CIOs face: modernizing legacy platforms without destabilizing the business. Blaze applies generative AI-driven orchestration to analyze existing systems; extract embedded logic; and improve continuity between requirements, development, testing, and deployment. This approach accelerates modernization while improving release predictability and reducing the risk that transformation initiatives introduce new operational fragility.



Managed technology services further support this capability architecture by stabilizing operational environments. By shifting defined run activities into predictable service models, organizations free internal teams to focus on higher-value work such as innovation, platform development, and AI integration.

When these capabilities operate together, modernization becomes a structured progression rather than a collection of disconnected projects. Organizations gain the ability to modernize legacy systems, scale AI capabilities, and maintain operational stability simultaneously.

Building the foundation for enterprise AI that works

The journey through the AI landscape demands more than just technological acumen; it requires strategic foresight, financial discipline, and decisive action from CIOs. The challenges of flat budgets, intense ROI pressure, and fragmented legacy systems are significant, but they are not insurmountable.

The CIOs who are succeeding in the AI era aren't those deploying the most tools. They are the ones strengthening the conditions that allow technology to function coherently across the enterprise.

Instead of just deploying tools, successful CIOs create the conditions for technology to thrive by establishing strategic clarity, building compelling financial narratives, modernizing foundational architecture, embedding adoption into the operating model, and strengthening operational resilience.

By adopting a comprehensive approach that prioritizes a defensible IT strategy built on a robust operating model, redefines value measurement, cleans up the data and architectural foundation, invests deeply in people, and controls costs with rigor, CIOs can lead their organizations beyond incremental gains. They can close the pilot-to-production gap and achieve truly transformative, enterprise-level AI outcomes.

Organizations that achieve this balance don't necessarily run more AI experiments than their peers. They simply become the organizations where AI works—and where enterprise AI truly scales.

Contact us



David Muir

Advisory Managing Director
US Technology & Data Strategy
KPMG LLP
drmuir@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us in:  | kpmg.com

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation. | MGT-9377

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.