



**The new CIO reality:
Five pressures
quietly reshaping
how enterprise
AI gets done**

KPMG.com



Enterprise AI ambition is rising faster than the systems designed to support it

CIOs are more visible in the enterprise than they've been in decades. AI has moved technology leadership from the operational layer of the organization into the center of strategic conversations.

Boards now ask CIOs not only how systems function but how technology investment will shape competitive advantage, operational resilience, and regulatory exposure.

With increased visibility, however, has come a new set of constraints, and CIOs are less in control than they appear to be.

CIOs are expected to expand AI capabilities, modernize infrastructure, strengthen security posture, and reduce technical debt simultaneously. Yet the financial and organizational structures supporting technology have changed far less dramatically than the expectations placed on them.

Technology budgets remain largely flat in many organizations, while operational demands—from cloud infrastructure to cybersecurity resilience—continue to grow.

At the heart of the issue: Most organizations treat AI adoption as a tech initiative, while the real constraint is the operating model required to make AI durable.

The result is a structural mismatch. AI initiatives promise transformative outcomes, but the conditions required to deliver those outcomes—clean data foundations, modernized architectures, disciplined lifecycle management, and organizational adoption—are often treated as background work rather than strategic investment.

For CIOs, this dissonance creates a moment of unusual complexity. The pressure to deliver AI results has accelerated faster than the enterprise operating model required to sustain them.

Across industries, five interlocking pressures are shaping how CIOs navigate the AI era:

- A funding paradox that demands more capability with limited investment growth
- Technology infrastructure and data foundations that struggle to support enterprise-scale AI
- Talent shortages and challenges around workforce adoption that determine whether AI investments generate real value
- Governance ambiguity around who owns AI risk and accountability
- External economic, regulatory, and geopolitical forces that compress timelines and increase scrutiny

Although each pressure can be manageable in isolation, the real challenge is that they rarely appear separately. Instead, they compound—creating a leadership environment where technology strategy, operating model design, and organizational alignment must evolve together.

Understanding how these pressures interact is becoming one of the defining leadership challenges of the AI era.

CIO self-assessment: Are the organizational conditions for enterprise AI actually in place?

Many organizations evaluate AI readiness through technology capabilities alone. A more revealing test examines the surrounding operating environment.

Consider the following questions:

- Can your organization clearly explain where AI costs sit and where AI value appears in financial reporting?
- Do your most promising AI pilots have a defined path to production deployment across the enterprise?
- Is AI adoption measured through workflow change and productivity impact rather than tool deployment counts?
- Could you defend your AI governance model to regulators and auditors tomorrow?
- If external conditions shifted—new regulation, vendor disruption, cost pressure—would your architecture adapt without major redesign?

If several of these questions are difficult to answer confidently, the pressures described in this paper may already be shaping your AI strategy more than you realize.

Find out more about the KPMG IT Maturity Assessment >



1. The funding paradox: Delivering more capability without structural budget growth

The Challenge

Technology leaders today face a paradox that rarely arises explicitly in boardroom conversations. Organizations expect AI-driven transformation, yet the economic structure supporting IT overall has changed very little.

Most CIO budgets remain dominated by operational commitments: maintaining legacy systems, managing vendor contracts, supporting cloud infrastructure, and ensuring security and compliance. These obligations leave limited funds to address foundational work that would enable AI at scale. Such work—architecture rationalization, data integration, platform consolidation, and lifecycle modernization—often compete directly with operational priorities that cannot easily be deferred.

At the same time, expectations surrounding AI investment continue to rise. Business leaders anticipate productivity gains, improved decision quality, and faster product innovation. And boards are increasingly asking how AI will translate into

competitive advantage and when will they see the ROI from AI that they were implicitly promised. Yet the enabling work behind those outcomes is rarely discussed in any significant way.

This leaves CIOs in an apparent no-win situation. The investments required to support enterprise AI frequently appear indirect when compared with initiatives that promise visible results within a fiscal year.

Why it Matters

But the funding challenge is not simply about budget size. The deeper issue lies in how organizations attribute and measure value.

Boards often focus on the wrong metrics—such as headcount reduction or immediate impact on profit and loss. The focus can move CIOs to overpromise on the board's desired metrics and fail to fully explain the true value AI generates, which at this point can be difficult to measure. AI benefits are real, but incremental, and often appear within business units—in faster analysis cycles,

improved customer interactions, or operational efficiency gains—not in the broad transformations that boards were promised. The infrastructure costs that enable those outcomes accumulate within the IT budget. Over time, the financial narrative becomes fragmented: value appears distributed across the enterprise while costs remain centralized. This disconnect weakens the CIO's position in funding discussions. Technology leaders are asked to justify infrastructure spending without a complete view of the value it enables.

Organizations that navigate this tension have successfully expanded the conversation beyond traditional ROI metrics. Instead of focusing narrowly on cost reduction, they demonstrate how AI initiatives influence cycle time, operational risk, forecast accuracy, and decision quality—outcomes that align more closely with the strategic objectives boards actually care about.

What leading CIOs are beginning to do differently

Some CIOs are reframing technology investment conversations in ways that restore coherence between cost and value. Rather than defending infrastructure spending as a necessary expense, they connect technology capability directly to business outcomes and enterprise risk reduction.

This shift often includes establishing clearer frameworks for measuring AI value, sequencing modernization initiatives so early improvements demonstrate tangible progress, and aligning



technology investment more explicitly with enterprise objectives. When the conversation moves from defending spend to explaining capability, technology investment becomes easier for boards and finance leaders to evaluate.

2. Infrastructure and data readiness: The hidden barrier to AI scale

The Challenge

Many organizations have successfully launched AI pilots. Far fewer have scaled those capabilities reliably across the enterprise.

The reason rarely lies in the AI model itself. More often, it reflects the condition of the surrounding technology environment. Most enterprise infrastructures weren't designed for AI workloads that involve dynamic computing demand, continuous data movement, and orchestration across interconnected systems.

As with any new technology, AI initiatives developed as isolated pilot projects rather than enterprise capabilities. Each new initiative required custom integration work, security validation, and operational oversight. When problems occurred—as they inevitably do—diagnosing them was a time-consuming process due to the complexity of interdependencies, resulting in sustained risks to production operations. While these pilots produced a patchwork of “pockets of success,” they also uncovered critical barriers within IT systems to a cohesive AI capability and enterprise-wide deployment.

Why it Matters

Data challenges intensify this problem, as AI is only as good as the data they run on. AI systems expose weaknesses that traditional analytics tools could often tolerate—namely unclear data ownership, inconsistent lineage, fragmented access controls, and poorly maintained knowledge repositories.

Data issues aren't new—but AI has given the old data quality conversation new urgency. When these problems surface through AI outputs rather than internal reports, they become visible immediately to business leaders.

Infrastructure fragmentation also affects financial predictability. Cloud cost volatility becomes difficult to explain when architectural standards are inconsistent and integration governance is weak. Without clear visibility into what drives infrastructure spending, CIOs struggle to justify the architectural investments that would reduce that volatility over time.

Organizations that successfully scale AI recognize that infrastructure readiness is not a background activity. It is the foundation for every AI initiative that follows.

The pilot-to-production gap in enterprise AI

Many CIOs recognize a recurring pattern:

1. A pilot demonstrates strong AI capability within a controlled environment.
2. The organization attempts to scale the solution.
3. Integration complexity, governance requirements, and data readiness slow progress dramatically.

The pilot did not fail. The enterprise environment simply was not prepared to support it.

Closing this gap requires treating architecture and data foundations as strategic investments rather than background maintenance work.

What leading CIOs are beginning to do differently

Technology leaders who successfully scale AI often begin by reframing infrastructure discussions. Instead of positioning architectural improvements—like data governance, integration architecture, and platform consolidation—as technical upgrades, they connect them directly to the AI use cases business leaders want to pursue.



These CIOs typically focus on:

- Creating reusable integration and security patterns that reduce per-project overhead
- Establishing shared platforms that provide consistent governance and observability
- Sequencing architecture improvements so they unlock multiple AI initiatives simultaneously

This approach gradually transforms fragmented environments into systems capable of supporting enterprise-scale AI.

Presented with that perspective, business leaders are more inclined to back efforts like data governance, integration architecture, and platform consolidation—especially when it's clear these initiatives help overcome challenges in achieving targeted AI applications.

3. Talent, culture, and adoption: Where AI value actually materializes

The Challenge

Deploying AI tools does not guarantee meaningful impact. Many organizations discover that employee adoption—not technology—is the decisive factor in whether AI investments succeed.

The pattern of workforce reaction to AI has been remarkably consistent. Executive sponsors advocate for AI strategically but rarely use the tools themselves. Middle managers worry about the implications for team structures. Process owners continue to rely on familiar workflows rather than experiment with new approaches.



Too often, the result is partial adoption. Tools exist, but behavior remains unchanged, and AI usage is low and uneven.

Why it Matters

Without measurable adoption, CIOs cannot demonstrate the value necessary to justify continued investment. When adoption stalls, skepticism grows, and the funding cycle becomes increasingly difficult.

Organizations also face the challenge of hiring specialized roles such as machine learning engineers or AI architects, which are expensive and competitive. But the larger issue is enabling widespread AI fluency across the workforce. Achieving high AI usage is a training and culture challenge as much as a hiring one.

Organizations that succeed recognize that AI adoption is fundamentally dependent on an operating model change rather than a technology rollout.

What leading CIOs are beginning to do differently

Rather than positioning IT as the sole driver of AI capability, leading CIOs are building operating models that encourage AI use. Technology teams encourage business units to experiment with AI tools, within shared platforms with established governance guardrails.

This model often includes:

- Executive education programs that help senior leaders experience AI directly
- Embedded AI technology capability within most or all business functions to encourage workplace-wide adoption and discourage the mindset that AI is only for “techies”
- Clear governance frameworks that allow workplace-wide AI experimentation without creating risk gaps
- A structured change-management process can mitigate risks and remove barriers to adoption, including phased deployment, active listening to employee concerns, continual training, and revised performance metrics to incentivize AI use

When adoption becomes a responsibility shared among all employees at all levels, AI value spreads faster and technology leadership gains credibility as a strategic enabler.



4. Governance and accountability: Defining ownership for AI risk

The Challenge

AI systems increasingly influence decisions that affect customers, operations, and regulatory compliance. Yet many organizations are still determining who is accountable when those systems produce unintended outcomes. This uncertainty is becoming more noticeable to boards as they confront these challenges.

Questions that once seemed theoretical are now operational concerns. Who owns responsibility when an AI output introduces bias, exposes sensitive information, or influences decisions that cannot withstand regulatory review? Traditional IT governance frameworks were not designed for the speed and complexity of modern AI systems. Model risk, data lineage, algorithmic accountability, and the auditability of AI-enabled workflows require a different set of controls. And right now, most organizations are building those controls reactively, after the fact.

Why it Matters

Boards are beginning to view AI risk similarly to cybersecurity risk—as a material business issue with direct implications for regulatory exposure and brand reputation. The question is no longer whether to govern AI, but whether the governance framework is credible—defensible to auditors, regulators, and the board itself.

CIOs who cannot demonstrate credible governance frameworks may find their authority challenged precisely when board attention is highest. Conversely, those who establish clear oversight models gain influence in strategic conversations.





What leading CIOs are beginning to do differently

Forward-looking organizations are building governance frameworks that combine technical oversight with business accountability. These frameworks often address:

- Model risk and validation
- Data lineage and auditability
- Vendor dependencies within AI supply chains
- Incident response protocols for AI-driven systems

By treating AI governance as a strategic capability rather than a compliance exercise, CIOs reposition technology leadership as a central component of enterprise risk management.



5. The external pressure cooker: When technology strategy meets global instability

The Challenge

The pressures facing CIOs described above rarely exist in a vacuum; they are often amplified by external forces. Economic uncertainty shortens planning horizons and intensifies budget scrutiny. Regulatory frameworks for AI continue to emerge unevenly across jurisdictions. Geopolitical dynamics introduce new considerations around vendor geography, data sovereignty, and technology supply chains.

Each of these forces pushes technology strategy in different directions. Economic pressure encourages faster proof of value. Regulatory pressure requires slower, more deliberate integration of AI into operational workflows. Geopolitical pressure demands transparency into infrastructure dependencies and vendor ecosystems.

These competing demands create a difficult balancing act. In environments with unknown dependencies, incomplete test coverage, and brittle point-to-point architecture, every AI release is a risk calculation.

CIOs facing external volatility are required to manage challenges beyond their direct influence, often while operating on infrastructures that allow little margin for error. Therefore, technology leaders must maintain long-term strategic direction while adapting to conditions that can shift dramatically from quarter to quarter.

Why it Matters

When these pressures collide with fragile technology environments, their impact multiplies. A deployment failure in a fragile environment not only leads to technical complications, but also invites regulatory oversight, disrupts scheduled activities, and requires immediate attention from the entire IT organization.

Organizations operating on resilient architectures can absorb disruption without derailing strategic initiatives. Those operating on brittle foundations often experience cascading consequences.

A single deployment failure can trigger regulatory scrutiny, stall modernization programs, and redirect engineering capacity toward emergency response.

External pressure therefore exposes underlying structural weaknesses where every shock hits harder and recovery takes longer.

What leading CIOs are beginning to do differently

The organizations navigating this environment most effectively share a common trait: they design for resilience before disruption occurs.

Organizations that absorb disruption without derailing their programs typically have made a few deliberate choices:

- Building architecture flexibility that allows systems to evolve safely
- Reducing vendor concentration risk through diversified ecosystems
- Establishing clear data residency and sovereignty policies
- Strengthening operational resilience through improved testing and dependency visibility



These efforts may not generate immediate attention, but they significantly expand the organization's capacity to adapt when external conditions shift. CIOs who proactively present operational resilience as a fundamental requirement for advancing AI initiatives, rather than merely a maintenance issue, often engage in more productive discussions than those prompted by system outages.

Final thoughts: Navigating the five pressures together

The pressures shaping the CIO role today are not new individually. Budget scrutiny, infrastructure complexity, adoption challenges, governance requirements, and external uncertainty have existed for years.

What is new is how tightly they now interact.

A CIO who resolves the funding narrative but ignores infrastructure readiness will see AI initiatives stall in production. One who modernizes architecture but neglects adoption will struggle to demonstrate value. Governance gaps can introduce regulatory exposure even when technology initiatives succeed technically.

Technology leadership in the AI era therefore requires a broader perspective. CIOs must address the full system—funding structures, infrastructure readiness, organizational adoption, governance frameworks, and external resilience.

Boards are not expecting perfection. They are looking for leaders who understand the complexity of the environment and can guide the organization through it with clarity and discipline.

Those CIOs are not simply implementing technology. They are redesigning how technology operates within the enterprise.



Contact us

Talk to us about how we can help you achieve your enterprise AI objectives.



David Muir

Managing Director
U.S. IT Strategy Lead
KPMG LLP

+1 512 320 5200

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:



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.

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

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.

