kyndryl

The technology reckoning

What a year of AI, quantum, cyber risk and infrastructure constraints reveals about the next era of enterprise leadership

By Kris Lovejoy

Global Head of Strategy, Kyndryl

Introduction

In 2025, I argued that the next phase of technology risk would not be defined by a single breakthrough, platform or threat actor. It would be defined by convergence: AI becoming operational infrastructure, quantum shifting from scientific possibility to strategic liability, data centers colliding with energy and water constraints, and digital systems moving deeper into the physical world.

Twelve months later, that thesis has largely held. Technology strategy can no longer be separated from capital strategy, geopolitical strategy, workforce strategy, sustainability strategy or operational resilience. The companies performing best are those best managing this convergence.

With this new era well underway, the next technology cycle will be shaped by compounding second-order effects, as enterprises discover that the hardest part of technology transformation is control.

Part I: The 2025 predictions, scored

Prediction is useful only if it is later subjected to evidence. Too often, technology forecasting rewards confident provocation rather than accountable analysis. The more disciplined exercise is to ask three questions: Did the predicted trend occur? Did it occur on the anticipated timeline? And did the original framing identify the right strategic implication?

On that basis, my 2025 predictions earn a mixed but instructive score. I was broadly right that enterprise technology risk would become more systemic, more physical and more entangled with policy and economics. I was also right that AI would accelerate both productivity and threat velocity. Where the forecast was less precise was in the mechanism. Some risks did not arrive as discrete shocks. They arrived as operating constraints: budget overruns, permitting delays, infrastructure scarcity, governance gaps and new forms of dependency.

That distinction matters. The scorecard below is not a victory lap or a mea culpa. It is an attempt to separate signal from noise before making a new set of predictions.

The scorecard

I would score the 2025 forecast as follows: three predictions were directionally right and ahead of schedule, one remains plausible but not yet proven, and one was correct in theme but too narrow in framing. The important lesson is that the center of gravity has shifted from technology adoption to technology consequence.

01 The autonomous adversary — Score: Mostly right, faster than expected

- **What I predicted:** AI-driven cyberattacks would move toward end-to-end autonomy, with meaningful enterprise impact by 2027.
- **What happened:** The forecast was broadly right, and the timeline compressed. Public reporting on AI-orchestrated cyber activity showed that adversaries were no longer using AI merely for advice, phishing copy or code generation. They were beginning to use agentic capabilities across reconnaissance, vulnerability discovery, exploitation, lateral movement, credential harvesting and data analysis.
- **What I would change:** The better frame is not "fully autonomous attack" versus "human-led attack." The more relevant risk is the progressive transfer of tactical execution from humans to machines. Even partial autonomy changes the economics of cyber offense by increasing speed, scale, and persistence.
- **Leadership implication:** Security operating models built around human-speed detection, ticket triage and manual response will continue to lose ground. Enterprises need defensive automation, adversary simulation against AI-enabled attack chains, and controls that assume machine-speed abuse of identity, APIs and software pipelines.

02 The disrupted cybersecurity workforce — Score: Plausible, not yet proven

- **What I predicted:** Economic dislocation in the technology workforce would push some highly skilled professionals toward cybercrime and increase the sophistication of ransomware and cybercriminal operations.
- **What happened:** The broader labor-market premise remains plausible, but the causal chain is difficult to prove. Cybercriminal operations have become more professional, more automated and more productized. However, attributing that shift directly to displaced technology workers would require stronger evidence than is currently available.
- **What I would change:** I would recast this prediction around capability diffusion rather than workforce displacement. The more important development is that advanced tools, cloud-native practices, AI coding assistants and commercial-grade operating models are becoming available to a wider range of actors.
- **Leadership implication:** Companies should update threat models for a world in which adversaries can access high-quality automation, software engineering practices and customer-service-like criminal infrastructure without needing elite internal talent.

03 The geopolitical splinternet — Score: Right, but understated

- **What I predicted:** Data localization and digital sovereignty would become tools of economic leverage and geopolitical contestation.
- **What happened:** This prediction was right, although the mechanism has been broader than anticipated. Sovereignty is no longer just a legal requirement or data residency question. It is becoming an architecture principle, a procurement criterion, a market-access condition and a negotiating lever in cross-border technology relationships.
- **What I would change:** I would broaden the frame from the “splinternet” to “sovereign operating environments.” Enterprises are not merely navigating fragmented internets; they are operating across jurisdictions that increasingly expect local control over data, compute, identity, cloud operations and AI governance.
- **Leadership implication:** Geopolitical analysis must move upstream into architecture decisions. Data placement, model training, support access, encryption, cloud control planes and operational continuity can no longer be treated as implementation details.

04 The cryptographic quantum threat — Score: Right direction, long fuse

- **What I predicted:** The risk of "harvest now, decrypt later" would force enterprises to treat quantum readiness as a strategic security issue well before a cryptographically relevant quantum computer is publicly demonstrated.
- **What happened:** That prediction remains directionally right. The public breakthrough has not arrived, but the migration challenge has become clearer. Post-quantum cryptography is no longer an academic discussion. It is a long-horizon enterprise transition involving cryptographic inventories, vendor dependencies, embedded systems, identity infrastructure, regulated data and long-lived secrets.
- **What I would change:** The prediction should focus less on the date of quantum decryption and more on the duration of migration. The risk is not that quantum suddenly breaks encryption tomorrow. The risk is that organizations wait too long to understand where cryptography lives across the enterprise.
- **Leadership implication:** Cryptographic agility should be treated as a funded transformation program, not a security patch. Boards should expect management to maintain a crypto inventory, prioritize high-value data classes, and build migration plans tied to vendor roadmaps and regulatory expectations.

05 Energy grid fragility — Score: Right, and now a board issue

- **What I predicted:** Data center growth, AI workloads and aging energy infrastructure would collide, creating availability and resilience risks for enterprises dependent on hyperscale capacity.
- **What happened:** This prediction was right — and the trend accelerated. Power availability, grid interconnection timelines, water use, cooling requirements, transformer shortages and local opposition have become constraints on the AI buildout. The bottleneck has shifted from ambition and capital to physical infrastructure.
- **What I would change:** I would describe this less as "grid fragility" and more as "the physicalization of digital strategy." Digital growth now depends on electricity, water, land, permitting, supply chains and community acceptance.
- **Leadership implication:** Cloud and AI strategy must incorporate physical resource exposure. Vendor diligence should include power sourcing, water usage, regional permitting risk, grid dependency, sustainability commitments and contingency plans for capacity rationing.

Part II: The technology trends that will define the next 12 months

The next 12 months will be less about discovering entirely new technologies than about absorbing the consequences of technologies already deployed. AI will move from pilots into operating models. Quantum will remain a future capability but become a present governance obligation. Data centers will become political and environmental flashpoints. Robotics and autonomous systems will turn cyber risk into physical risk. And enterprises will discover that the hardest part of technology transformation is not invention. It is control.

Trend 1: AI moves from experimentation to economic discipline

The first phase of generative AI was funded as exploration. Organizations authorized pilots, productivity experiments, coding assistants and proof-of-concept deployments with the expectation that long-term strategic value would justify early investment.

That phase is ending. AI is now becoming a run-rate expense, an operating dependency and, in some cases, a source of budget volatility.

The economic problem is not simply that tokens cost money. It is that agentic AI changes the nature of consumption. A single user request can trigger planning, tool calls, retrieval, retries, model-to-model exchanges, code execution and validation loops. The organization may see one business interaction while the underlying system performs dozens or hundreds of model operations.

In addition to token consumption, emerging conversations about open-weight models introduce another strategic consideration, raising questions about whether every AI engagement must run through a "metered" service. The conversation shifts from "How many tokens will this require?" to "What's the most efficient model for this workload, based on the outcomes I need?"

This will force a new discipline: AI FinOps.

Traditional cloud FinOps was designed around infrastructure units that were relatively visible and attributable. AI consumption is more fluid. It can be embedded in applications, triggered by agents, hidden in workflows or generated by employees using tools outside formal procurement channels. The unit economics may improve even as total spending rises.

Prediction: Within the next 12 months, at least one major enterprise will publicly reset its AI program because consumption costs outpaced measurable value creation. The companies that avoid this outcome will establish real-time token visibility, multi-model-routing discipline, semantic caching and explicit value metrics before scaling agentic workflows broadly.

Trend 2: The Enterprise creates an ungoverned AI estate

As employees continue to experiment with AI, the greatest risk is that experimentation will create a new estate of agents, copilots, low-code automations, scripts, plug-ins, model connections and AI-enabled applications that the enterprise cannot see or govern.

This estate differs from shadow IT in one critical respect: many AI-enabled systems do not merely store information or automate a static workflow. They can reason across context, retrieve data, call APIs, use credentials, initiate transactions, generate code and interact with other agents. They are full actors within the enterprise environment. Yet in recent research from Kyndryl, just 27% of executives told us they had a central registry and monitoring of all AI systems operating in their environments.

The governance gap will become apparent when leaders ask basic questions that should be easy to answer: What agents exist? Who owns them? What data can they access? What actions can they take? Which models do they depend on? Are they still needed? What happens if they behave unexpectedly?

Traditional software estate	Emerging AI agent estate
Applications are usually registered, owned and deployed through formal processes.	Agents and automations may be created rapidly by business teams, developers or vendors.
Behavior is largely deterministic and bounded by predefined workflows.	Behavior may vary based on prompts, context, tools, model outputs and agent-to-agent interactions.
Cost and risk are typically visible through infrastructure, licenses and application owners.	Cost and risk may be hidden in token consumption, API calls, credentials and autonomous execution paths.

Prediction: Within the next 12 months, a significant enterprise incident will be traced not to a flawed foundation model, but to an unregistered or poorly governed AI-enabled workflow. That incident will accelerate demand for an enterprise AI control plane: a governance layer that discovers, registers, secures, monitors, and retires agents and AI-enabled applications across the organization.

Trend 3: Digital growth becomes constrained by physical resources

The AI boom has made an old truth newly visible: digital infrastructure is physical infrastructure. Models require data centers. Data centers require power, water, land, cooling, transformers, substations, fiber, skilled trades, permits and community acceptance. The constraint on AI adoption is increasingly not imagination or capital. It is the physical layer beneath the cloud.

This will change enterprise technology procurement. Buyers will begin asking cloud and AI providers questions that once belonged to utilities, environmental teams and local governments. Where is the capacity located? What is the source of power? How exposed is the region to grid congestion, water stress, permitting delay or community opposition? Can the provider sustain commitments during localized rationing or infrastructure delays?

Enterprises will increasingly need to treat infrastructure legitimacy as a strategic input. The social license to operate will matter. Communities that see data centers as engines of tax revenue may support expansion. Communities that experience them as drivers of higher electricity prices, water stress or land-use conflict may resist them.

Prediction: Within the next 12 months, enterprise cloud and AI sourcing decisions will begin to include explicit assessments of power availability, water usage effectiveness, regional permitting risk and provide social license to operate. These criteria will move from sustainability reports into procurement scorecards.

Trend 4: Quantum becomes a governance problem before it becomes a compute advantage

Quantum computing will remain one of the most important technology trends of the coming year. The near-term issue is governance.

Organizations must prepare for a future in which cryptographic assumptions change, optimization methods improve and industry-specific quantum use cases begin moving from research partnerships toward early workflow integration.

The mistake would be to wait for a dramatic quantum breakthrough before acting. Cryptographic migration is slow because cryptography is everywhere: in applications, devices, certificates, identity systems, network protocols, third-party products, embedded systems, backups, archives and long-lived regulated datasets. Many organizations do not yet know where their vulnerable cryptography sits.

At the same time, quantum will increasingly matter for competitive strategy. Industries such as pharmaceuticals, materials science, financial services, logistics, energy and advanced manufacturing will continue exploring hybrid quantum-classical approaches. The first-mover advantage may not come from owning a cryptographically relevant quantum computer but from owning the data, workflows, talent and intellectual property around high-value use cases.

Prediction: Within the next 12 months, leading enterprises will separate quantum readiness into two programs: one focused on post-quantum cryptographic migration and one focused on business-value discovery through quantum and hybrid quantum-classical use cases. The former will be governed as risk reduction; the latter as option creation.

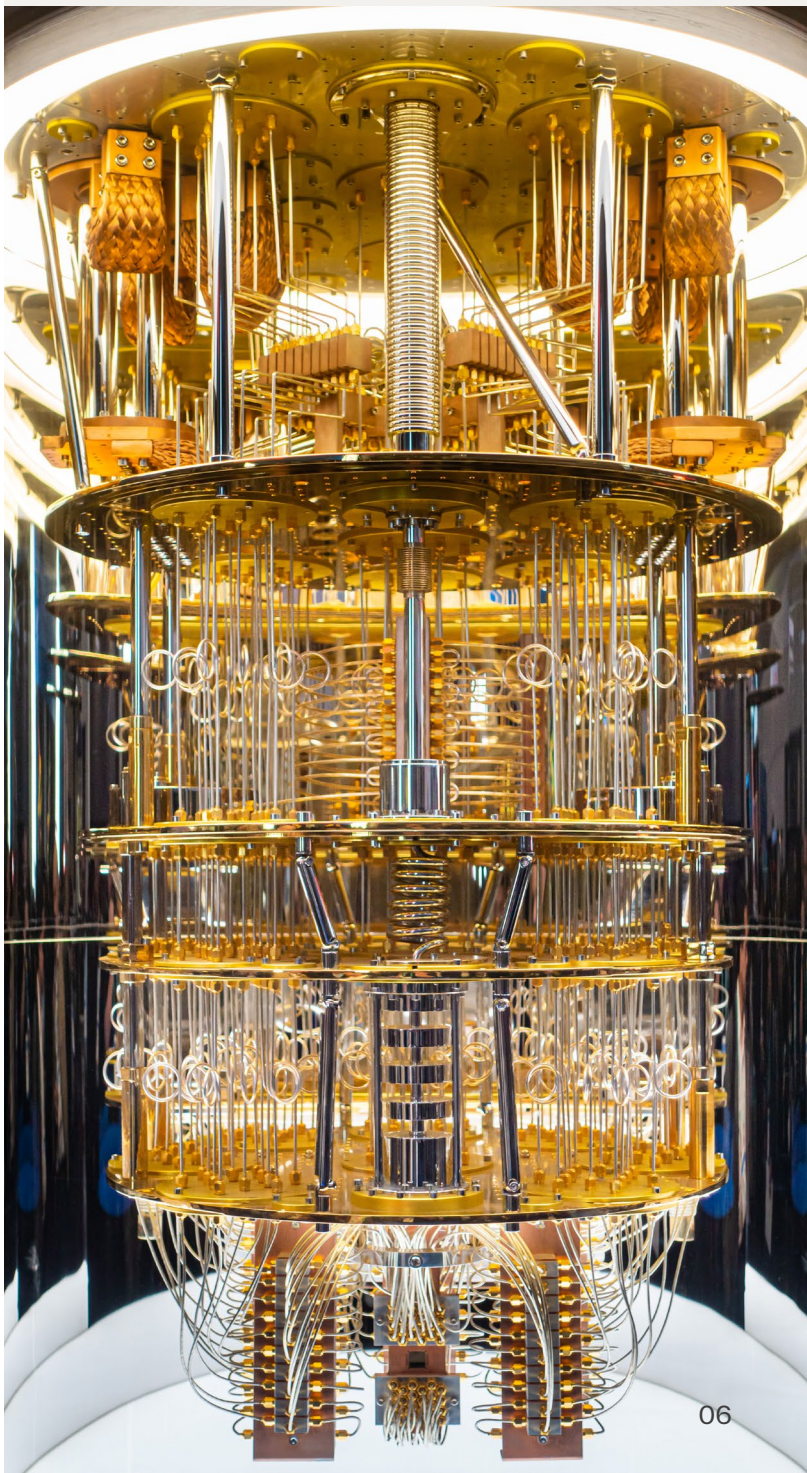

Part III: The executive agenda

The lesson from the last year is not that leaders should be more cautious about technology. It is that they should be more disciplined. The companies that benefit most from AI, quantum, automation and cloud infrastructure will be those that build the management systems to govern them. Five actions should be on the executive agenda.

1. Treat AI as an operating model, not a tool portfolio

AI governance should extend beyond model approval and acceptable-use policies. It should define how work changes, who owns agentic workflows, how value is measured, what controls are enforced at runtime and how exceptions are handled.

- Establish clear ownership for AI-enabled workflows, not just AI tools.
- Measure value at the business-process level rather than at the pilot level.
- Require runtime controls for data access, tool use, model selection and escalation.

2. Build AI FinOps before scaling agentic workflows

The economics of AI should be visible before adoption accelerates. Finance, technology, procurement, risk and business owners need shared instrumentation for token consumption, multi-model routing, value realization and exception management.

- Track token consumption by workflow, owner, model and business outcome.
- Use multi-model routing so routine tasks do not default to the most expensive frontier models.
- Apply semantic caching, context discipline and usage limits to reduce waste.

3. Create an enterprise AI control plane

Enterprises need a governed inventory of agents, copilots, AI-enabled applications, model connections, tool permissions, data entitlements and life cycle status. Without that inventory, leaders cannot manage risk, cost or accountability.

- Register every agent and AI-enabled workflow with an accountable owner.
- Classify each agent by data access, action authority, model dependency and business criticality.
- Monitor for anomalous behavior, policy violations, excessive token consumption and orphaned permissions.

4. Integrate physical resource exposure into digital strategy

Boards and executive teams should expect technology strategies to include the physical assets that underpin them. Cloud regions, AI capacity, data center partners and edge deployments should be evaluated for power, water, permitting and community risk.

- Include regional energy and water exposure in cloud and AI vendor diligence.
- Ask providers for contingency plans related to capacity rationing or permitting delay.
- Align sustainability commitments with the infrastructure required to support AI growth.

5. Fund quantum readiness as a multi-year program

Quantum readiness should be managed before urgency becomes crisis. The first step is knowing where cryptographic exposure exists and where quantum-enabled business advantage may eventually matter.

- Maintain a current inventory of public-key cryptography across applications, infrastructure, products and third-party dependencies.
- Prioritize migration for long-lived, regulated or strategically sensitive data.
- Identify industry-specific quantum use cases where early learning could create defensible advantage.

The leaders who succeed will not be those who predict every trend perfectly. They will be those who build institutions capable of learning quickly, governing thoughtfully, investing deliberately and adapting when the forecast changes. In that sense, the technology agenda for the coming year is about whether management systems can evolve fast enough to govern technological change.

kyndryl

© Copyright Kyndryl Inc. 2026. All rights reserved.

Kyndryl is a trademark or registered trademark of Kyndryl, Inc. in the United States and/or other countries. Other product and service names may be trademarks of Kyndryl, Inc. or other companies.

This document and the information contained herein are provided solely for informational and Kyndryl marketing purposes and should not be relied upon as advice or a recommendation.