

A woman with dark hair is shown in profile, looking out of a circular window. The window reveals a vibrant, digital representation of the Earth, overlaid with a network of glowing blue and yellow lines that suggest global connectivity or data flow. The scene is set against a dark, starry background, giving it a futuristic, space-themed appearance. The woman's hands are clasped together in front of her.

2026 Frontiers Report

AI as the frontier multiplier

July 2026

Table of contents

PART ONE / INTRODUCTION

- 03 Opening Letter
- 04 AI as the frontier multiplier
- 06 Frontiers

PART TWO / THE FRONTIERS

- 08 Path to AGI
- 14 Physical AI
- 21 Energy & Compute
- 28 Quantum & NextGen Computing
- 40 Future of Trust
- 45 Space
- 52 Areas We're Watching

PART THREE / CONCLUSION

- 55 Conclusion
- Appendix: Glossary
- Appendix: Methodology
- Appendix: Sources

Leadership in an interconnected world

The pace of change is accelerating, driven by advances in technology, shifts in global systems, and increasing pressure on infrastructure and resources. Last year's frontiers report focused on convergence, the idea that multiple forces are evolving simultaneously and intersecting in new ways.

Those intersections are now more visible than ever. What once appeared as parallel trends now operate in combination and progress in one domain directly influences outcomes in others. Developments in computing affect energy demand. Advances in artificial intelligence (AI) expand the capabilities of physical systems. Breakthroughs in biological engineering are beginning to reshape manufacturing and supply chains. Activity in space is extending the reach of communications, sensing, and data.

AI is linking these developments as a catalyst and accelerator.

More and more, artificial reasoning systems are embedded in the operations, infrastructure, and decision-making processes that organizations rely on to function and compete. As capabilities improve and deployment expands, AI affects how systems learn, make decisions, and coordinate activity.

The first frontier in this report, Path to AGI, examines this shift directly. It focuses on how enterprises are integrating intelligence into decision-making, governance, and operations. In practice, the path towards advanced machine intelligence is unfolding not only in research labs but inside organizations as they reimagine how work gets done.

This year's report examines several frontiers where these changes are most visible. AI is present across all of them and, in many cases, is accelerating how quickly change spreads between them. In each domain, new capabilities are emerging that promise to change our world in novel and profound ways.

Predicting a single future is not the purpose of this report. Instead, it provides a structured view of the forces shaping the landscape and the interactions that connect them. Understanding these dynamics can help leaders make more informed decisions about investment, partnerships, and operating models.

The goal is to help leaders understand not only what is changing, but how these changes connect and how those relationships influence risk, resilience, and competitive positioning.

In conversations with executives, several questions appear frequently:

- How should organizations invest when technological change affects multiple domains at once?
- How can leaders scale new capabilities while maintaining governance and control?
- How can companies strengthen resilience while continuing to pursue growth?

These questions sit at the intersection of technology, strategy, and leadership.

The 2026 Frontiers Report is intended to support leaders facing an unprecedented period of dynamic change. It highlights where cross-domain effects are intensifying, how AI is influencing those developments, and where strategic decisions today may shape outcomes in the years ahead.

We hope the report serves as a useful reference as leaders navigate the choices ahead.



Brian Fields
KPMG Enterprise Innovation Leader



Stephanie Kim
KPMG Enterprise Innovation



AI as the frontier multiplier

AI now plays a structural role across the domains shaping the global economy.



AI is becoming embedded within computation, infrastructure, industrial production, biological manufacturing, and geopolitical competition. As capabilities improve, progress in one frontier increasingly accelerates developments in others, creating a multiplier effect.

Improvements in reasoning models influence robotics and automation. Advances in computing architectures expand deployment scale. Growth in data center capacity reshapes energy demand and infrastructure planning. Biological design increasingly depends on computational discovery tools. The pattern is also visible across industrial environments, where advances in machine perception, planning, and control enable intelligent automation. In biology, AI-driven tools make modeling protein structures, simulating metabolic pathways, and designing biological systems increasingly feasible.

AI functions less as a standalone sector and more as a connective capability embedded within the economy. Understanding these interactions is critical for strategic decision-making, where investments within one frontier shape opportunities in another.

This report examines six frontiers where AI advances are reshaping technological development, capital investment, and competitive advantage across computing, industry, infrastructure, biology, and global markets.

Key takeaways



AI is multiplying across systems, not replacing them.



Energy and governance are the real multipliers or constraints.



The frontiers are interconnected, not independent.



Competitive advantage shifts to orchestration.

What changed from the 2025 report


The 2025 Frontiers Report focused on convergence.

Last year, the challenge was understanding convergence. This year, the challenge is understanding how that convergence is being shaped by intelligence embedded across enterprise operations, infrastructure, and markets.

As volatility compounds across the economic, political, and technological landscapes, organizations are striving to adapt faster, make decisions sooner, and execute change more effectively.

Capabilities are compounding.	Execution is moving closer to the real world.	Constraints are becoming central.
Advances in reasoning, planning, multimodal understanding, and autonomy are no longer progressing in isolation. They are reinforcing one another, accelerating timelines and compressing development cycles.	Capabilities that were previously confined to advisory or experimental settings are now operating within production environments. Decisions are increasingly translated directly into action across infrastructure, workflows, and physical systems.	Energy availability, infrastructure readiness, compute capacity, regulation, and geopolitical exposure are no longer secondary considerations. They are shaping how and where systems can scale.

AI AS THE FRONTIER MULTIPLIER



Shaping the next era Frontiers

How the chapters are organized

The report begins with **Path to AGI**, which examines how organizations are embedding intelligence into decision-making, governance, and operations.

Subsequent chapters explore how advances in AI influence physical operations, computing architectures, biological production, energy demand, space activity, and the architecture of trust.

Each chapter provides an overview of the frontier, examines recent developments, and highlights signals that may influence its trajectory.

How to read the frontier chapters

Readers can approach the frontier chapters independently, depending on their interests and priorities. Taken together, the chapters provide a broader perspective on how developments across multiple domains influence one another and shape the environment in which organizations operate.

Taking the frontier journey

Within each frontier, look for the footer action bar titled **Frontier Journey: Building Your Blueprint**. These blueprints are designed to help readers move from insight to action through a visual summary of each frontier.

Each blueprint highlights:

- The shift
- Signals under the radar
- Why this matters
- What leaders are doing
- From insight to action

Each frontier blueprint is available for download in the **Take Action** section of the corresponding frontier chapter.

Frontiers



Path to Artificial General Intelligence (AGI)

AI is progressing toward systems capable of reasoning, planning, and coordinating complex work across organizations. As these capabilities mature, the central challenge for enterprises is shifting from access to intelligence toward the governance, control, and operating structures required to deploy increasingly autonomous systems safely and effectively.



Physical AI

AI is not just advancing in software; it is increasingly working with machines, infrastructure, and operational environments that interact with the physical world. As sensing technologies, robotics, and AI capabilities progress, intelligent systems are beginning to perform work across manufacturing, logistics, transportation, and other sectors of the real economy.



Energy & Compute

The expansion of large-scale computing is increasing demand for electricity, data center capacity, and supporting infrastructure. As AI workloads grow, the intersection of energy systems and computing infrastructure is becoming a strategic issue for governments, technology providers, and enterprise leaders.



Quantum and NextGen Computing

New paradigms and technologies are expanding the limits of classical computing. Quantum systems, neuromorphic chips, photonic processors, and specialized AI hardware are advancing rapidly. The future of computing is likely to involve hybrid environments in which organizations use multiple architectures together to solve different classes of problems.



Future of Trust

As advanced technologies become embedded across digital and physical environments, mechanisms that establish trust, verification, and accountability are becoming increasingly important. Blockchain systems, identity frameworks, AI verification tools, and governance models are emerging as critical components of enterprise architecture.



Space

Space is increasingly becoming an operational layer of global infrastructure supporting communications, sensing, navigation, and data systems. Expanding satellite networks and declining launch costs are enabling new commercial activity in orbit while influencing industries and supply chains on Earth.

Contributing Authors



Angelina Budko
Associate Director, KPMG
Enterprise Innovation



Carrie Young
Manager, KPMG
Enterprise Innovation



Priya Roy
Manager, KPMG
Enterprise Innovation





Path to AGI

When intelligence stops being scarce

Artificial General Intelligence (AGI) is usually framed as a race against the calendar. Will it arrive in five years, ten years, or longer? For two centuries, every institution we have built has assumed that thinking is expensive. You acquire it by hiring people, training them for years, buying the firms that already have it, or waiting for experience to accumulate. Cognitive capacity has been the hardest resource to scale and the slowest to move. Org charts, budgets, governance, career ladders, the very shape of the modern firm, are all downstream of that one constraint. Remove it, even partially, and the assumptions built on top of it start to give way.

AGI is often used to describe systems that can reason, learn, adapt, are polymath, transfer knowledge and operate effectively across a wide range of unfamiliar domains without task specific retraining. That definition is useful, but it buries the lead. But for enterprise leaders, the more important point is not whether a system satisfies a precise definition. It is that frontier AI systems are moving along that trajectory, becoming more capable, more general, and more useful across increasingly complex work.

The more important enterprise implication is this: the label matters less than the direction of travel. Enterprises do not need to wait for a formal declaration of AGI to feel the consequences. Reasoning, planning, memory, tool use, learning, machine teaching, and autonomous long-horizon execution are improving quickly and converging into systems that can take on more ambiguous, cross-functional work. METR (Model Evaluation and Threat Research)'s research, for example, suggests that the length of software tasks frontier agents can complete autonomously has doubled roughly every seven months over the past six years, with signs the trend may have accelerated since 2024.¹ Whether we call that AGI, frontier AI, or simply more capable AI, intelligence is becoming more elastic, and organizations need to learn how to direct, govern, and benefit from it.

The significance, then, is not whether AGI arrives on a precise date or meets a perfect definition. It is that intelligence is moving from a scarce organizational resource toward something that can be scaled, aimed, and governed as an economic unit.

The economics of scalable intelligence

Organizations have traditionally expanded cognitive capacity by hiring people, developing expertise, bringing in specialist firms, or investing in training. Intelligence has therefore been expensive, unevenly distributed, and difficult to scale.

As intelligence becomes more scalable, that equation starts to change. Organizations are no longer limited only by how many experts they employ. The limiting factor becomes how effectively they define problems, provide context, coordinate work, reimagine the work itself, govern decisions, and convert reasoning into outcomes. This does not diminish the importance of human expertise. Quite the opposite. Human intelligence increasingly defines objectives, exercises judgment, establishes values, manages tradeoffs, and evaluates outcomes. Scalable intelligence expands the organization's ability to reason, explore alternatives, synthesize knowledge, and learn across the enterprise.



Steve Chase
Principal, KPMG
Global Head of AI &
Digital Innovation



Swami Chandrasekaran
Principal, KPMG
Head of AI and Data Labs



The economic implications extend well beyond productivity. Organizations may rethink how they innovate, allocate capital, organize work, preserve institutional knowledge, and compete. Just as cloud computing changed assumptions about infrastructure, AGI or scalable intelligence may change assumptions about how cognitive work is created, distributed, and managed. This shift is already visible in enterprise behavior: according to KPMG's Q2 2026 Pulse Survey, the share of organizations orchestrating multiple agents across workflows has doubled since the previous quarter (9% to 18%),¹¹ and research from the Stanford Digital Economy Lab suggests that productivity differences increasingly reflect how effectively firms integrate data and redesign workflows around AI.¹² The binding constraint stops being whether you have enough cognitive capacity and becomes how well you aim it.

Capabilities that change the enterprise

The significance of AGI is not that it introduces entirely new capabilities. Early versions of many of these capabilities are already visible in today's frontier models and are converging. What changes is their breadth, depth, reliability, adaptability, and applicability. Capabilities that appear intermittently today may become routine, require less human guidance, and extend across domains, industries, and situations.

1

The first is generalization.

Can it reason through a problem it has not been explicitly trained or taught to solve?

AGI can apply existing knowledge to unfamiliar problems without requiring a predefined workflow or task-specific training. Many consequential enterprise problems have little precedent. During due diligence, for example, a company may encounter a target with an unusual business model, fragmented data, shifting regulation, and limited precedent. A more general system could reason across market structure, customer behavior, operational dependencies, contractual exposure, and value creation potential to form a view where no standard deal template applies. The value is not retrieving precedent. It is reasoning when precedent is limited.

2

The second is cross-domain reasoning.

Can it reason through a problem it has not been explicitly trained or taught to solve?

AGI can combine knowledge that traditionally sits in separate disciplines. A pharmaceutical company could explore new therapeutic hypotheses by integrating molecular biology, immunology, manufacturing science, epidemiology, genomics, behavioral health, and health economics. A global bank could reason across risk, tax, treasury, compliance, legal, operations, technology, and product strategy. The value is not faster analysis within a single function. It is the ability to identify relationships, tradeoffs, and pathways that are difficult to see from any one discipline alone.

3

A third capability is transferable learning.

Can it reuse insight gained in one context to improve performance in another?

Knowledge and insights from manufacturing around flow, capacity, variability, and resource coordination could improve hospital systems operations. Lessons from financial risk could inform supply-chain resilience or cyber threat detection, because they share underlying patterns of uncertainty, interdependence, and risk. Patterns learned from hundreds of acquisitions could help improve large transformation programs by identifying recurring indicators of integration risk, organizational readiness, and value leakage. This is materially different from automation. Automation repeats a known process. Transferable learning allows insight to move across the enterprise and compound.

4

Fourth, AGI also introduces long-horizon reasoning.

Can it reuse insight gained in one context to improve performance in another?

Rather than optimizing isolated activities, AGI can evaluate choices over months or years and revise recommendations as conditions change. Planning becomes adaptive rather than episodic. A more capable system could continuously assess geopolitical shifts, demographic trends, regulatory change, technology disruption, capital markets, energy constraints, supplier fragility, and competitive moves to help guide portfolio strategy, investment decisions, or multi-year transformation programs as conditions evolve.

5

Fifth, adaptive intelligence.

Can it learn, reason about its own reasoning, and continuously refine its understanding of how the world works?

Beyond solving individual problems, AGI may be able to build and update internal models of complex systems, recognize uncertainty in its own conclusions, challenge assumptions, incorporate new evidence, and improve from experience. During M&A due diligence, for example, it could move beyond document review to build a continuously evolving view of how a company creates value, where risks emerge, and which assumptions matter most to the investment thesis. As new information arrives, the system updates its understanding, identifies gaps in evidence, and adapts its recommendations. The shift is from static expertise to intelligence that learns, refines, and improves over time.

6

Sixth, invention and discovery.

Can it generate new hypotheses, designs, products, molecules, or operating models from what it has learned?

Discovery and invention are the outcomes of creative synthesis at scale. A system that can generalize beyond precedent, reason across disciplines, transfer knowledge, learn from experience, and adapt its understanding of complex systems can surface possibilities that may not emerge within a single team or function. In healthcare, research, treatment planning, care delivery, reimbursement, capacity management, and patient outcomes could begin to operate as a connected learning system. New discoveries could influence care pathways sooner, and clinical outcomes could reshape future treatment decisions. The opportunity is not simply a more efficient healthcare organization. It is a system capable of improving itself.

Collectively, these capabilities change the role intelligence plays inside the enterprise. Instead of helping people perform individual tasks, intelligent systems increasingly participate in how organizations learn, adapt, discover, plan, and make decisions. As these capabilities mature, intelligence itself becomes a more scalable resource, reshaping how enterprises create value, manage risk, and compete.



Enterprises have a decisive role to play in the evolution of AGI and superintelligence.
It cannot be defined solely by a handful of labs at the frontier.

Swami Chandrasekaran Principal, KPMG, Head of AI and Data Labs

What leaders should be preparing for

The precise moment at which a system qualifies as AGI is less important than the direction of capability improvement. Organizations do not need to wait for a formal AGI milestone to experience the consequences. As AI systems become more capable at reasoning, learning, planning, discovery, and adaptation, enterprises will need to rethink assumptions that were built around intelligence being scarce, expensive, and difficult to scale.

The basic shape looks like this:

- Workforce planning becomes cognitive capacity planning:**
 Workforce planning has traditionally determined how much cognitive work an organization could perform. As intelligence becomes increasingly scalable, leaders may need to think about cognitive capacity alongside financial capital, physical assets and human talent. This is critical because for centuries, organizations have built human expertise through an apprenticeship model. As scalable intelligence automates this layer of entry-level work, a trend accelerated by recursive improvement techniques seen in systems like Google DeepMind's AlphaEvolve², the traditional talent pipeline is at risk of breaking. The strategic question shifts from "How many people do we need?" to "How much reasoning capacity do we need, where should it be applied and how should it be governed?" Some enterprises are already redefining the nature of work; Shopify's CEO instructed teams that before requesting additional headcount, they must demonstrate why AI cannot perform the task in question.¹⁵
- Knowledge becomes more than documentation:**
 Organizations have always depended on more than what is written down. Tacit expertise, informal networks, accumulated judgment, customer nuance, project history and context that exists only in people's heads often matter as much as formal documentation. As AI systems become more capable, enterprises will increasingly find ways to preserve, externalize and reason over both explicit and tacit knowledge. Documents, policies, engineering standards, contracts and research will remain important, but they will no longer exist only to be read. They become part of the living context that informs decisions as they are made.
- Institutional memory becomes infrastructure:**
 Most organizations lose valuable context every day. Decisions disappear into meetings. Rationale gets buried in emails. Teams repeat work because prior learning is hard to find. Knowledge leaves when people leave. In a world of scalable intelligence, that becomes a structural disadvantage. Living memory systems that preserve decisions, assumptions, relationships, evidence, lessons learned, and organizational knowledge may become as important as today's data platforms. Organizations that remember better will reason better, which is why firms like JPMorgan Chase and Goldman Sachs are centralizing AI deployment to build enterprise-wide "cognitive infrastructure."¹⁴
- Technical debt becomes a strategic constraint:**
 As intelligence becomes more scalable, technical debt becomes more visible and more expensive. Legacy systems, fragmented data, brittle integrations, undocumented business rules, and aging controls increasingly limit an organization's ability to deploy AI safely and at scale. More capable AI systems and models may also expose vulnerabilities, control gaps, and process weaknesses that were previously masked by human effort. The competitive advantage will increasingly belong to organizations whose technology, data, and trust/control environments are ready for scalable intelligence, not those forced to navigate around legacy constraints.
- Sophisticated workspaces evolve:**
 Email, presentations, meetings, task lists and chat are unlikely to disappear. But they may be complemented by persistent cognitive workspaces where people and intelligent systems explore ideas, evaluate alternatives, retain context and build on previous work instead of repeatedly recreating it. This is not simply a better collaboration tool. It is a different surface for enterprise work, one where human and machine intelligence operate against shared context over time.

- **Governance becomes continuous. Always ON:**

Existing governance models were designed around human decision velocity. They rely on periodic reviews, approvals, audits and controls. As increasingly capable systems participate in more consequential decisions, governance must operate inside the flow of work. Policies, guardrails, verification and assurance become embedded within operational processes rather than applied after the fact. Identity, provenance and observability become foundational. The regulatory environment is already formalizing this shift; ISO/IEC 42001 is emerging as a management framework for AI governance,³ the NIST AI Risk Management Framework is being used to guide risk management,⁴ and the EU AI Act's high-risk obligations take effect in August 2026, as part of a phased implementation.⁵ Organizations need confidence in where knowledge came from, how conclusions were reached, what influenced a recommendation and whether systems remained aligned with enterprise objectives throughout the reasoning process.

- **Identity becomes the new, non-negotiable control plane:**

Because an agent's claim on compliance checks cannot be taken at face value, identity becomes the foundational layer for governance. Public reporting highlights the urgency: In Okta's 2025 executive survey, 91% of organizations reported already using AI agents, while only 10% had a strategy for managing non-human identities.⁶ Microsoft's February 2026 Cyber Pulse report said more than 80% of Fortune 500 companies are already using active AI agents,⁷ yet Cloud Security Alliance and Strata Identity report that only 18% of organizations are highly confident their current IAM systems can manage agent identities.⁸

- **Cognitive security becomes a board-level issue:**

Protecting information is no longer sufficient. Enterprises will also need to protect the integrity of objectives, reasoning, memory, accumulated knowledge and decision pathways. This security challenge exists at three levels: First, at the operational level, as "shadow agents" proliferate; KPMG's 2025 CEO Outlook finds that 86% of CEOs expect AI agents to function as embedded team members within the next year.⁹ Second, at the legal level, where liability for autonomous decisions remains institutionally unresolved.¹⁰ Finally, at the macroeconomic level, as regulators act on systemic risks; a 2024 report from the Financial Stability Board highlights concentration across key AI providers as a potential source of cascading risk.¹³

- **Context becomes the moat:**

Access to powerful models is unlikely to remain exclusive. But enterprise context will be harder to copy. Customer relationships, operating experience, institutional knowledge, decisions, proprietary data, organizational memory, trusted principles, judgment and accountability become more important as intelligence becomes more widely available. Intelligence may become abundant. Context does not. As access to powerful models becomes widespread, advantage shifts to trusted relationships, institutional knowledge, judgment, accountability, and the operating context in which decisions are made.

- **The leadership challenge:**

None of these shifts require the arrival of a universally accepted definition of AGI. They emerge as AI systems become increasingly capable and increasingly embedded in how organizations create knowledge, make decisions, preserve context, allocate capital, and manage risk. The strategic challenge for leaders is therefore not predicting when AGI arrives. It is determining how quickly these capabilities are improving, where they create competitive advantage, and how their organizations will adapt as intelligence becomes increasingly scalable.



AI is not simply a technology program. It represents
a fundamental rewrite of how business operates.

Steve Chase, Principal, KPMG Global Head of AI & Digital Innovation

A simple shift to remember

Enterprise today	Enterprise with scalable intelligence
Workforce determines cognitive capacity	Cognitive capacity becomes elastic
Knowledge is captured in documents	Knowledge becomes continuously accessible and extensible
Governance is periodic	Governance becomes continuous
Institutional memory decays	Institutional memory compounds
Expertise is scarce	Context is scarce
Intelligence is expensive	Judgment is expensive

Conclusion: when intelligence stops being scarce

Each technological era changes what is abundant and, in doing so, changes what is precious. Industrialization made physical labor more scalable and elevated design, capital, and management. Computing made calculation abundant and elevated software. The internet made information access abundant and elevated networks, platforms, and attention. As intelligence becomes more scalable, the prize moves to what abundance cannot easily touch: judgment, trust, context, governance, accountability, leadership, and human relationships.


The organizations that succeed will not necessarily be those with access to the most capable models. They will be those that know how to combine scalable intelligence with deep institutional knowledge, disciplined governance, resilient cognitive security, and clear human accountability.

The future of AGI remains uncertain. Its timing, technical path, and ultimate capabilities will continue to be debated. But the strategic direction is already visible. AI systems are becoming more capable, more general, and more embedded in how organizations reason, learn, decide, and act.

The question facing leaders is therefore not whether AGI has officially arrived. It is whether their organizations are prepared for a world in which intelligence becomes increasingly scalable, becomes the economic unit, and no longer serves as the constraint around which the enterprise was designed.

FRONTIER JOURNEY
Building your Blueprint

Download the **Path to AGI Blueprint**
to help navigate your journey.





Physical AI

The AI race has left the screen

The biggest AI story is no longer about the models running on screens; it's about the intelligence now acting in the physical world. This is not a distant future. For enterprise leaders, the provocative question has shifted from whether AI will reshape their operations to if their operating model can survive when AI begins to act, not just advise. When intelligence is embedded into machines that move products, inspect assets, and navigate dynamic environments, the consequences become immediate. This shift marks the arrival of Physical AI, and it is moving along an exponential curve.

Earlier waves of AI changed how people write, search, analyze, and decide. Physical AI changes how work gets done in the world. When intelligence is embedded into machines that move products, inspect assets, operate equipment, and navigate dynamic environments, the consequences become immediate: productivity can rise, but so can safety, liability, and operational risk. The provocative question for leaders is no longer whether AI will reshape the enterprise. It is whether their operating model is ready for AI to act, not just advise.

Why now?

The forces accelerating Physical AI are capital, capability, operating pressure and advances in models:

- **Capital is scaling deployment:** After declining in 2022–2023, Physical AI investment rebounded in 2024–2025. Venture funding in robotics and physical AI exceeded \$16B in 2025.¹ In October 2025, SoftBank agreed to acquire ABB's robotics business in a deal valued at about \$5.4B.²

SILICON VALLEY PHYSICAL AI SPEND

In USD billions



AI as the Frontier Multiplier

AI now empowers physical systems to learn faster and adapt in complex and unpredictable real-world environments. This exponentially expands the uses for automation and robotics across all major sectors. For companies, managing this transition strategically is essential, as AI moves from the digital screen into core operations. These capabilities could create a decisive competitive advantage for those that can bridge intelligent decision-making with AI-fueled physical execution.



Brian Higgins

Principal, KPMG

US Sector Leader for Industrial Manufacturing



Jono Anderson

Principal, KPMG

Industrial Manufacturing



Greg Corlis

Principal, KPMG

Tech Data Strategy



Chris Corteen

Senior Director, KPMG

AI and Data Labs



- **Operating pressure is rising:** Physical AI is gaining traction not only because the technology is improving, but because operations are under strain. Labor shortages, scarce skills, and supply chain disruption are pushing companies to automate targeted tasks now, especially where reliability and continuity matter most.
- **Rise of open world foundation models:** Physical AI is becoming easier to build because companies no longer have to start from scratch. Shared AI models, simulation tools, and synthetic data let them train and test systems faster and lower development costs. But this does not eliminate a core constraint: high-quality real-world training data for robots remains hard and expensive to collect, especially for edge cases, physical interaction, and dynamic environments. That is why simulation and synthetic data matter so much—they help narrow, but not eliminate, the data bottleneck.

Why it matters?

Physical AI raises the stakes of AI adoption. In software, mistakes can slow work or reduce quality. In physical systems, they can affect safety, uptime, asset performance, reputation, and financial results. This makes Physical AI a matter of operating model, governance, and capital planning. Strategically, this means competitive advantage will rapidly shift from basic process efficiency to systemic autonomy. As this transition occurs, operational risk calculations must be completely rewritten, recognizing that a software failure can now trigger a direct physical-world catastrophe. Consequently, capital allocation will begin to bifurcate, heavily favoring organizations that integrate adaptable, intelligent systems over those that continue to invest in rigid, traditional machinery.

The shift

As Physical AI moves from experimentation to scaled deployment, enterprise system design is evolving in parallel.

Physical AI should not be conflated with fully general-purpose humanoid robots. The near-term market is broader and more practical: autonomous mobile robots, industrial systems, warehouse automation, inspection platforms, and domain-specific embodied agents operating in constrained environments. Humanoid robots may become an important long-term frontier, but most enterprise value in the next several years is more likely to come from narrower systems that perform reliably in defined tasks and settings.

Here's how we know the change is happening:

- **Intelligent task orchestration:** AI systems are moving beyond single-purpose automation toward handling multiple actions within defined environments. That is not the same as fully general-purpose robotics. Progress is strongest where tasks, environments, and failure modes can be bounded. Models now translate language and perception into real-world execution. The 2025 Stanford AI Index highlights progress in multimodal and embodied AI⁶, while Google DeepMind's RT-2 model shows how language commands can be translated into robotic action.³
- **Edge-based autonomy:** Decision-making is moving to the "edge," placing computing power directly where it is needed. The global edge AI market is projected to reach \$386 billion by 2034, reflecting rising demand for real-time processing in physical environments.⁴ That demand is driving adoption of the specialized, low-power processors designed to run inference locally without relying on cloud connectivity.
- **Software-defined robotics:** The real breakthrough in Physical AI is increasingly in software, not hardware alone. New model architectures are making it possible for machines to interpret broader instructions, understand context, and translate perception into action. That does not eliminate the importance of hardware, but it shifts more of the differentiation into the software layer: models, data, orchestration, and continuous improvement through updates rather than redesign.
- **Simulation-first design:** Companies are using virtual models to test physical systems before deployment. The digital twin market is expected to surpass \$140 billion by 2031⁵, as industries from manufacturing to energy use these virtual replicas to optimize designs, predict maintenance needs, and reduce risks before committing capital.

The boundary between digital decision-making and physical execution is narrowing. A new ecosystem of interconnected, autonomous systems is beginning to take shape across industries.

The next wave: how industries will fundamentally change:

- **For industrial manufacturing:** For Industrial Manufacturing: The next wave isn't just more automation; it's the adaptive factory. Picture a production line that reconfigures itself in real-time based on supply chain disruptions, with robots learning new tasks overnight from simulations. This transforms manufacturing from a rigid process into a dynamic, intelligent organism.
- **For logistics and fulfillment:** The next wave isn't just warehouse robots; it's the autonomous supply chain. Imagine a logistics network where autonomous trucks, drones, and sorting systems operate in a continuous, coordinated flow from port to final delivery, managed by an AI that anticipates demand and reroutes shipments before bottlenecks even form.
- **For infrastructure and energy:** The next wave isn't just remote monitoring; it's self-maintaining infrastructure. Think of drones and robotic crawlers that don't just inspect pipelines and power grids but perform repairs on the spot, guided by AI that predicts failures weeks in advance based on subtle thermal and acoustic signals.

Signals landscape

4.7M

Operational industrial robots in 2024

International Federation of Robotics 2025

82%

Manufacturers cite a lack of AI-ready skills as their top challenge

Manufacturing Leadership Council 2025

\$250 B+

Global edge computing spend 2025 forecast

IDC 2025 Guide

100 M+

Fully autonomous miles driven

Waymo Safety Report 2025

~1 to 10

Robot-to-worker ratio, South Korea

IFR Robot Density 2025

700 K+

Projected (2028) global robot installation

IFR 2025 Outlook

1 M+

Robots operating in Amazon facilities

Amazon 2025 reporting



With AI, the total addressable market was centered on knowledge workers. **Physical AI expands that scope across industrial systems and operational labor.**

Brian Higgins, Principal, KPMG, US Sector Leader for Industrial Manufacturing

Signals under the radar

Early signals suggest that Physical AI is beginning to reshape both operations and the institutional systems around them. Regulation, insurance, and governance are evolving in response to autonomous systems operating in the real world.

1

Regulation is beginning to shape system design upstream

Regulatory frameworks are starting to influence how autonomous systems are built, not just how they are evaluated. Requirements around safety, explainability, and auditability are embedded earlier in development cycles.

INDICATORS:

- High-risk AI classifications influencing product design decisions and compliance requirements early in development⁶
- Early alignment of engineering standards with regulatory frameworks
- Procurement requirements specifying auditability and documentation

2

Liability is shifting from humans to systems

As autonomous systems take on decision-making, responsibility is beginning to move from operators to algorithms, vendors, and system owners. Insurance models and regulatory frameworks are starting to adapt, though standards remain fragmented.

INDICATORS:

- Emerging insurance models for autonomous systems⁷
- Growth in robotics and industrial automation underwriting
- Emergence of algorithm-specific liability and accountability frameworks

3

Verification is moving from periodic review to embedded system behavior

As autonomous systems become more prevalent, the need for ongoing monitoring and validation increases. NIST's AI Risk Management Framework reinforces this shift toward lifecycle governance.⁸

INDICATORS:

- Integration of monitoring into operational systems
- Simulation and testing embedded into deployment workflows
- Procurement requirements for continuous performance validation



Intelligence in motion demands intelligence at the edge: the ability to think in milliseconds, without waiting for the cloud, without hesitation.

Edge AI is the nervous system that makes Physical AI alive.

Greg Corlis *Principal, KPMG, Tech Data Strategy*

As we witness AI evolve from generating insights to driving real-world actions, the risks become more severe. Failures can now cause operational disruption, equipment damage, or in the worst cases, harm to people. Yet adoption of Physical AI is accelerating faster than the safeguards needed to ensure safety for workers and the public. Without purpose-built safety frameworks, organizations face a growing gap between capability and control. Key considerations will include ensuring safe human-machine interaction, validating AI-driven actions in dynamic environments, and establishing accountability for autonomous decisions.

4

Operational data is entering financial and compliance systems

Data from physical systems is beginning to feed directly into reporting, audit, and risk management. Digital twins and real-time telemetry are enabling greater visibility into asset performance and system behavior.

INDICATORS:

- Adoption of digital twins for lifecycle management and reporting
- Integration of operational data into financial and compliance systems
- Use of real-time asset telemetry in audit functions

5

Cross-border deployment is becoming more complex

Physical AI systems must navigate fragmented regulatory environments across regions. Differences in safety standards, certification, and data rules are increasing deployment complexity. Physical AI also depends on globally contested supply chains for semiconductors, sensors, compute infrastructure, and industrial components. As localization pressures, sovereign technology strategies, and divergent standards increase, deployment decisions become not only a question of compliance, but of resilience, sourcing, and regional operating model design.

INDICATORS:

- Divergent high-risk classifications across jurisdictions
- Separate certification regimes for robotics, mobility, and industrial systems
- Data localization requirements affecting edge deployment
- Region-specific safety validation and inspection protocol

“

Nearly everything is going to become impacted by robotics, whether it's a car, our offices, homes.

Anything that is arduous and repeatable will benefit heavily from the Physical AI.

Jono Anderson, Principal, KPMG, Industrial Manufacturing

How enterprises are responding

Physical AI is shifting from a technology initiative to an operating model decision. Organizations are rethinking value chains, governance, and infrastructure.

Scaling investment and capability

Organizations are committing capital and leadership attention to long-term Physical AI capabilities.

- Global industrial robot installations neared 575,000 in 2025 and are projected to top 700,000 per year by 2028.⁹ Manufacturers are accelerating smart-factory automation, signaling a permanent shift away from rigid, single-purpose assembly lines toward highly flexible, dynamically reconfigurable production environments.
- Tesla is shifting its manufacturing focus to include humanoid robots, converting existing EV production lines at its Fremont facility to produce its Optimus robot while breaking ground on a new, larger robotics plant at Giga Texas.¹⁰ This represents a strategic bet that the future of labor relies on general-purpose machines capable of learning any human task, fundamentally altering industrial economics.
- Logistics and retailers are aggressively scaling robotics to boost throughput and resilience. Amazon has deployed over 1 million robots across its fulfillment network, reflecting the scale at which physical automation is being integrated into core operations¹¹ and establishing a competitive moat built on physical speed and resilience that will be nearly impossible for rivals to replicate quickly.
- Technology firms are also investing in embodied AI and robotics capabilities. Meta is advancing research in robotic manipulation and telepresence through its AI research programs¹², laying the groundwork for a future where digital AI agents can seamlessly perceive and manipulate the physical world.
- Hyundai has invested over \$400M to establish the Boston Dynamics AI Institute, focusing on foundational advances in robotics and physical AI¹³, to solve the most complex hardware-software integration challenges and accelerate the commercialization of fully autonomous machines.

Moving from pilots to deployment

Enterprises are shifting from isolated pilots to real operating environments.

- Edge computing investment continues to accelerate, supporting localized execution in industrial and logistics environments.
- Digital twins are increasingly used to validate workflow changes before physical rollout.
- Organizations are scaling across multiple sites to build operational confidence before full rollout.

Building enterprise infrastructure

While major technical barriers remain for generalizable robotics, many near-term enterprise deployments will be limited less by model capability than by integration, monitoring, safety, and governance at scale.

- Companies are establishing centralized governance models to manage risk and compliance
- Real-time monitoring and evaluation systems are being built to track performance and detect issues
- Middleware and integration layers are being standardized to enable repeatable deployment across environments



Physical AI is more than robotics. It is the integration of AI in physical systems that can sense, navigate, and interact with the real world.

Chris Corteen, Senior Director, KPMG, AI and Data Labs

From insight to action

Physical AI is moving into early deployment across industries. Organizations are beginning to introduce autonomous systems into core operations, where they directly affect performance, safety, and reliability.

This introduces new requirements. Deployment depends on integrating these systems into existing operations, managing risk in real time, and scaling them across environments.

LEADERS SHOULD FOCUS ON TWO QUESTIONS

Operational readiness for autonomy

Can the organization safely deploy and manage autonomous systems across critical operations?

Exposure to physical-system failure

Does the organization understand the safety, financial, and liability risks if autonomous systems fail, and are mitigation measures in place?

A Structured Path Forward: Now, Near, Next


These questions translate into a set of decisions that unfold over time. The path from early deployment to scaled operations requires different actions at each stage.

Now	Near	Next
- Fund 2–3 high-value deployments tied to core operations (e.g., warehouse automation, predictive maintenance, autonomous inspection) - Establish executive ownership for Physical AI across operations, technology, and risk (not isolated in innovation teams) - Build minimum viable governance, safety, cybersecurity, and trust controls to enable deployment without slowing execution	- Scale deployments across sites and business units where ROI and reliability are proven - Standardize platforms and integration layers to avoid fragmented systems and vendor lock-in - Shift operating model from human-led to human-supervised systems in targeted workflows	- Redesign operations around autonomous execution (factories, logistics, infrastructure running with minimal human intervention) - Reallocate workforce and capital toward system oversight, exception handling, and optimization - Own or control critical system layers (data, orchestration, safety, decision logic) that define competitive advantage
DECISION: Where do we place our first capital bets, and who owns the outcome?	DECISION: Which systems become core to how we operate, and what do we standardize vs. keep flexible?	DECISION: What do we need to control to compete, and where are we willing to rely on external platforms?

Physical AI extends intelligence into the systems that produce and maintain the physical economy. As a result, decisions across technology, operations, safety, cyber security, trust and risk are becoming more tightly linked. The challenge is not access to technology but the ability to integrate it into existing operations. This requires clear accountability, disciplined risk management and systems that can operate reliably under changing conditions. Organizations that approach Physical AI as a set of isolated deployments are likely to encounter limits to scale. Those that align it with operating model design and capital allocation will be better positioned to adopt it in a sustained way.

FRONTIER JOURNEY
Building your Blueprint

Download the **Physical AI Frontier Blueprint** to help navigate your journey.





Energy & Compute

The New Constraint—Power, Chips, and AI Scale

For decades, companies could expand computing without worrying about power. Electricity was abundant, predictable, and grew quietly alongside digital demand. That reality is now changing as a structural shift emerges between two foundational forces: computing power, the processing capacity behind cloud services and artificial intelligence, and energy, the electricity required to run it.

Historically, compute expanded quickly while energy systems grew more slowly and remained largely outside strategic decision-making. Today the two are becoming closely connected, and the availability of power is shaping where and how quickly data centers and digital systems infrastructure can grow.

Why now?

Artificial intelligence systems use computing power in ways previous digital services did not. Training large AI models requires thousands of powerful chips running at the same time in a single location, which creates major spikes in electricity demand at those sites. Once deployed, these systems run continuously, increasing the amount of electricity required across many data centers. The International Energy Agency projects global data center electricity demand will roughly double by 2030, reaching around 945 TWh in its base case.¹

At the same time, the systems that produce and deliver electricity expand slowly. Building new power plants, grid connections, and transmission lines requires lengthy permitting, financing, and construction.² Technology cycles are measured in months, while energy infrastructure often takes years to expand.³ As a result, digital growth is increasingly shaped by physical limits that cannot be quickly scaled.

Why it matters

The rapid growth of AI-driven computing is colliding with constrained power systems, turning electricity from a predictable operating cost into a potential business risk. Access to power, energy prices, and the pace of infrastructure development are becoming board-level considerations. This is most visible for hyperscale companies building data centers, but it increasingly affects any enterprise whose growth depends on digital infrastructure.

Running AI at scale requires more than electricity. It depends on a chain of specialized, hard-to-get hardware — and while the chips themselves are relatively available, the high-bandwidth memory (HBM) those chips require to function and the manufacturing process that assembles them into something deployable are not. These shortages are so severe that major manufacturers like Micron booked their entire 2026 HBM capacity by December 2025.⁴ Energy and hardware are the most visible constraints, but they are not the only ones. AI now depends on a stack of four interdependent layers, each of which can slow the entire chain.

AI as the Frontier Multiplier

This frontier is being shaped by physical and economic limits that are arriving sooner than expected. AI-driven demand for compute is stressing grids, data center infrastructure, cooling systems, and supply chains. Permitting, financing, and geopolitical access to critical inputs are shaping where capacity can be built and who can scale. These trends explain why energy and compute are no longer enabling layers. They are becoming binding constraints that influence strategy across every other frontier.



Brad Stansberry
Principal, KPMG
Energy and Utilities



Vijay Subramanyam
Principal, KPMG
Technology, Media
and Telecommunications



Reid Tucker
Principal, KPMG
Capital Projects Advisory



The compute constraint: from power outages to “token brownouts”

While the availability of electricity is the most visible constraint, a second, more subtle challenge is emerging: the resilience of access to compute itself. The discussion of operational resilience is shifting beyond securing power supply and toward securing AI inference capacity.


As AI becomes embedded in core business processes, from customer service to financial reporting, enterprises will require continuous, guaranteed access to the AI models that run them. However, as energy and hardware constraints intensify, the market for AI inference may become volatile. This introduces a new form of operational risk: the “Token Brownout.”

A “Token Brownout” is a scenario where an enterprise cannot access the AI inference capacity it needs, when it needs it. This may not be a complete outage but a degradation of service—higher latency, lower priority, or simply an inability to secure capacity for critical workloads. Unlike traditional cloud computing, Quality of Service (QoS) for large-scale AI inference is not yet a mature, guaranteed market. This creates a critical strategic question for leaders: not just if the data center has power, but if the business has priority access to the AI compute its operations now depend on.

Because these compute resources and hardware layers are so difficult to obtain, running AI remains expensive, and guaranteed access remains uncertain. To manage these costs and access risks, enterprises are strategically segmenting where workloads run—using cloud platforms for elasticity and frontier model access, while shifting to smaller, less expensive, or locally hosted models for predictable, high-volume inference workloads.

Figure 1. The AI scaling stack

Four layers of constraint now shape how quickly enterprises can deploy AI. Each layer depends on the one below it. **A breakdown at any layer constrains the layers above. AI deployment depends on all four.**



The shift

In the past, enterprises have expanded compute capacity, AI workloads, and data-driven operations without regard to energy constraints. Scaling and expanding digitally is now subject to issues such as grid shortfalls, rising electricity prices, and infrastructure build timelines that cannot match AI-driven demand growth.

Here's how we know the change is happening:

- **Grid stress:** Grid operators are signaling stress. PJM Interconnection's December 2025 capacity auction came up 6.5 gigawatts short of its reliability target for the 2027-28 delivery year. This indicates infrastructure expansion is failing to keep pace with load growth.⁵

- **Hyperscaler spending boom:** Capital flows confirm the mismatch. Market reporting suggests major tech platforms could spend roughly \$600 billion in 2026, up 70% from 2025,⁶ with a large share directed toward AI compute and next-generation data centers. In 2025, five of the largest hyperscalers issued \$121 billion in U.S. corporate bonds, versus an average of \$28 billion per year between 2020 and 2024.⁷ U.S. data center construction spending has risen sharply to roughly \$40 billion annually by mid-2025.⁸
- **Compute becomes a strategic constraint:** The global data center boom is not just about building more facilities; it's a race to secure strategic assets. Major technology firms are not only spending billions on construction but are also engaging in increasingly aggressive M&A and talent acquisition to control every layer of the AI stack, from chips to the developers who use them. Access to high-end compute is a vector of competitive advantage and a potential single point of failure.
- **Utility investment scaling:** Utilities are accelerating investment, but on longer timelines. Forty-three publicly traded utilities project \$1.1 trillion in capital investment from 2025–2029. Transmission and distribution investment is forecast to remain well above 2018–2020 levels.⁹
- **Rising energy costs:** Price signals reinforce stress. U.S. retail electricity prices have increased since 2022 across customer classes, and residential prices remained elevated through 2025.¹⁰ Electricity is no longer a stable operating input. It is becoming a structural boundary on enterprise growth, capital velocity, and margin durability.

Signals under the radar

The most important changes are unfolding beyond forecasts and spending levels. They are appearing in how digital systems are designed, how projects are financed, and how risk is evaluated. Together, these signals suggest energy constraints are beginning to shape how digital systems grow.

1

Space-Based Compute

In May 2026, The Wall Street Journal reported Google and SpaceX are currently in talks to launch data centers into orbit.¹¹ This underscores how seriously the industry is searching for alternatives to terrestrial power, land, and cooling constraints. While still early and unproven, the conversation itself is a signal: when infrastructure bottlenecks become strategic, firms begin looking beyond Earth-bound optimization toward entirely new compute architectures.

2

Compute-Side Efficiency Breakthroughs

Neuromorphic and photonic chips aim to deliver more computing power while using less electricity, shifting scaling from GPU volume to performance per watt. For example, in April 2025, researchers demonstrated the first large language model running on Intel's neuromorphic Loihi 2 chip, a potential step toward AI inference that draws a fraction of conventional power.¹²

Advanced cooling systems and power-aware workload scheduling are also improving energy utilization. This is evident in the work of Corintis and Microsoft, who developed microfluidic cooling channels etched directly into chips — with architecture bio-inspired by butterfly wing veins — that remove heat three times more effectively than conventional cold plates.¹³

3

Energy-Side Breakthrough Bets

Capital is shifting toward long-horizon, reliable energy sources aligned to AI-scale demand, suggesting companies want more direct control of energy supply. Power sourcing is becoming a strategic control point, influencing where infrastructure is built and who controls it to secure supply. Fusion firms raised about \$2.64 billion in the 12 months to July 2025.¹⁴ Nuclear startups also drew substantial funding in 2025. Meanwhile, the U.S. Department of Energy has committed up to \$800 million in cost-shared funding to support early deployment of advanced small modular reactors in the U.S.¹⁵

4

Expanding Risk Surface

As digital and power systems converge, industrial control systems that run energy infrastructure are becoming cybersecurity risks. Grid-level, digital, and financial risks are merging into a growing operational risk. Check Point research recorded an average of 1,162 cyberattacks per week on utilities through August 2024, up from 689 over the same period in 2023, a nearly 70% increase.¹⁶ In April 2026, a joint advisory from the FBI, CISA, NSA, DOE, and U.S. Cyber Command confirmed that Iranian-affiliated actors had actively disrupted U.S. energy, water, and government infrastructure by exploiting industrial control systems.¹⁷

5

Social and Regulatory Friction

Regulatory and social resistance is no longer a peripheral nuisance. It is becoming a hard constraint on AI infrastructure growth. Community opposition, court rulings, and permitting scrutiny can now stall, shrink, or stop large-load projects outright, turning siting risk into a direct threat to speed, cost, and capacity. Tucson's Project Blue faced major opposition over water and energy concerns.¹⁸ A judge voided rezoning tied to the Digital Gateway project near Gainesville, Virginia.¹⁹ Litigation also challenged industrial revenue bond approvals tied to Project Jupiter in Doña Ana County.²⁰ What begins as local backlash can quickly become a systemic brake on expansion.

What leaders are doing

Building Cross Sector Coordination

Leading enterprises now recognize that energy and compute constraints are inseparable—and that no single company or industry can solve them alone. This recognition is driving unprecedented cross-sector collaboration.

Industry coalitions are bridging the tech-power divide. EPRI's DCFlex initiative brings together hyperscalers and data center operators (Google, Microsoft, Meta, NVIDIA, Compass, QTS) with major utilities (Duke Energy, Southern Company) and grid operators (PJM Interconnection) to jointly tackle data center flexibility and grid integration challenges.²¹

Utilities are also formalizing "build together" partnerships with tech customers. Duke Energy has signed agreements with Amazon, Google, and Microsoft to co-develop new large-customer structures and clean energy approaches in the Carolinas.²²

SEEING THE SHIFT IN ACTION

The AI infrastructure buildout is forcing a fundamental shift from siloed planning to coordinated governance, with tech companies, utilities, and regulators now working as interdependent partners rather than separate stakeholders.

Securing Behind-the-Meter and Co-Located Generation

Leaders are moving upstream into energy development and direct power control:

- xAI installed dozens of gas-powered turbines to provide interim power while awaiting sufficient grid capacity at its Colossus data center in South Memphis.²³
- Amazon acquired Pine Gate Renewables' 1.2 GW Sunstone solar-plus-storage project in Oregon through bankruptcy proceedings.²⁴
- The Department of Energy formally urged FERC to establish national rules for large loads connecting directly to transmission.²⁵
- FERC ordered PJM to establish clearer tariff rules governing the co-location of large loads—including AI data centers—with on-site generation.²⁶

SEEING THE SHIFT IN ACTION

When power is a utility purchase, growth depends on external timelines and tariffs. As a controlled asset, expansion and margins become more predictable, and energy shifts from operating expense to strategic asset.

Redesigning the Energy-Compute Stack

To overcome energy constraints, industry leaders are investing in technologies that address both sides of the constraint: expanding reliable energy supply and decreasing the energy required for computation

On the supply side, firms like Google²⁷ and Microsoft²⁸ are investing heavily in next-generation solutions, including advanced nuclear and fusion energy through various power purchase agreements and strategic partnerships.

On the demand side, companies are innovating for greater efficiency. A notable example is the partnership between robotics firm SAFELOG and startup Opteran Technologies. Through this collaboration, they aim to demonstrate that physical navigation tasks can be performed with significantly less energy by replacing power-intensive navigation systems with biologically inspired AI.²⁹

SEEING THE SHIFT IN ACTION

These investments reflect a shift from incremental optimization to structural redesign of the compute–energy interface, with enterprises simultaneously advancing efficiency and generation to hedge scarcity while reducing energy needed for each AI task.

Building Operational Resilience

Energy independence is extending beyond hyperscale data centers into broader enterprise operations. H-E-B invested in microgrids and on-site generation to allow stores to operate independently of the grid during storms or peak demand events.³⁰

SEEING THE SHIFT IN ACTION

Organizations are building resilience into their power supply — investing in backup generation and on-site energy assets to keep critical operations running when there are disruptions.

Active Energy Management

Leading companies now treat electricity as a strategic input, not a fixed cost—managing procurement, timing, and efficiency to protect margins and enable growth. On procurement, Alphabet agreed to acquire clean energy developer Intersect for \$4.75 billion in December 2025 to directly control large-scale power supply for its data center expansion.³¹ On timing, Voltus delivered \$240 million in earnings and savings to commercial and industrial customers in 2025 by managing 8.1 GW of flexible load.³² On efficiency, TSMC employed dynamic optimization to reduce the peak power consumption of its EUV (extreme ultraviolet) lithography tools by 44%, a result first achieved in 2024 and deployed at scale in 2025.³³ As power availability and cost become constraints on growth, leading companies are taking direct control of their energy position.

SEEING THE SHIFT IN ACTION

Integrating energy sourcing, infrastructure timing, and efficiency into capital planning allows organizations to expand with greater predictability and more stable margins, turning power from a passive operating cost into a strategic asset.



We are at a clear inflection point where energy availability is becoming a meaningful constraint on enterprise technology strategy...in many regions, the limiting factor for AI deployment is not capital or compute, but speed to power. **Enterprises that do not factor grid capacity, permitting timelines, and infrastructure readiness into their AI strategies risk delayed deployment, higher costs, or stranded investments.**

Arshad Mansoor, President and CEO of EPRI

What if?

The uncertain future of AI and energy creates a strategic asymmetry. Acting early to build an energy strategy provides valuable options across multiple futures, while waiting leaves you exposed regardless of which scenario unfolds.

Energy investments are not binary bets; they are a resilient foundation.

- If AI's energy demand continues to grow, early action improves the likelihood of securing necessary capacity and infrastructure.
- If new energy breakthroughs arrive, established partnerships and expertise can be positioned to evaluate and deploy them more quickly.
- Even if GPU efficiency eliminates the constraint, the organizational capability built—strategic planning, energy partnerships—becomes a competitive advantage for any future capital-intensive project.

From insight to action

Energy signals are clearer, constraints are increasingly quantifiable, and market responses are already underway. The advantage now lies in execution. The question is no longer whether energy will shape enterprise growth, but where and how leaders act to stay ahead of the constraint. The following sections outline the priority domains where immediate action plans can reduce exposure, preserve flexibility, and protect expansion velocity.

Establish an enterprise energy baseline

Consolidate energy consumption, pricing exposure, and load concentration across the enterprise to identify where the organization is most vulnerable to supply disruptions or price volatility.

ACTION:

Build a consolidated enterprise-wide view of energy consumption, pricing exposure, and load concentration before scaling new AI workloads.

Design for flexibility

As demand patterns and regulatory conditions continue to shift, generation, contracting, and siting decisions should be structured to preserve flexibility. That flexibility is sustained by pacing capital deployment alongside verified demand growth and infrastructure readiness, rather than relying solely on long-range projections.

ACTION:

Sequence capital investment against proven demand signals and confirmed infrastructure milestones.

Integrate energy into scenario planning

Model energy demand across growth scenarios and embed load, pricing, and interconnection assumptions into capital decisions.

ACTION:

Approve expansion only with clear visibility into power availability, cost durability, and timing risk.

Anticipate regulatory and community scrutiny

Expect increased review of grid impact, emissions, water use, and rate allocation.

ACTION:

Integrate stakeholder engagement and permitting risk into expansion planning.

Secure your compute supply chain


Map your organization's dependency on external AI model providers and inference capacity. Diversify your portfolio of AI providers and explore contractual options for securing priority access and Quality of Service.

ACTION:

Begin treating AI inference as a critical supply chain input, subject to the same resilience and diversification strategies as physical components.

FRONTIER JOURNEY
Building your Blueprint

Download the **Energy & Compute Frontier Blueprint** to help navigate your journey.





Quantum & NextGen Computing

Enterprise computing architecture is expanding from traditional general-purpose architectures to specialized architectures designed for different business and technical needs.

Quantum

For years, enterprise technology strategy followed a familiar pattern. First, the technology matured. Then standards emerged. Then regulators, investors, vendors, and large companies moved at scale. Quantum has not followed that sequence. Institutional momentum arrived first — and in 2026, the technology began to catch up to it.

Quantum technologies are already delivering commercially relevant results, even as fault-tolerant, error-corrected systems continue to mature. Enterprises are no longer preparing for a future technology. They are responding to one that has begun to arrive.

In 2024, the U.S. National Institute of Standards and Technology finalized its first post-quantum encryption standards, signaling the start of migration planning across federal systems and regulated supply chains.¹ In 2026, public capital accelerated sharply. The U.S. government announced approximately \$2 billion in awards to quantum computing firms, taking equity stakes in select companies to anchor domestic capability.² The Department of Commerce and IBM announced America's first purpose-built quantum foundry, supported by a proposed \$1 billion CHIPS Act award,²² and PsiQuantum signed a \$100 million letter of intent with the Department of Commerce in the same period.³ Major economies have embedded quantum in national strategies and committed billions of dollars in public funding across computing, sensing, security, and workforce development.^{4,5,6,7}

Not every company needs a quantum computer, but every company with long-lived sensitive data, regulated systems, or long investment horizons needs a quantum position.

There are four things to pay attention to:

- **Post-quantum cryptography (PQC):** updating today's quantum-vulnerable encryption so sensitive data stays protected as quantum machines advance toward breaking the methods in wide use today.
- **Quantum computing:** identifying where this technology solves certain hard business and scientific problems faster or better than today's systems.
- **Quantum sensing:** identifying where quantum-based sensors improve how we measure, detect, or navigate in areas like logistics, infrastructure, healthcare, and defense.
- **Hybrid compute architecture:** planning for an environment where quantum systems work alongside traditional computing, AI, and high-performance machines rather than replacing them.

AI as the Frontier Multiplier

The Quantum frontier is being shaped by exponential advances in computing, driven by the need for faster processing of AI model outputs. Changes are being tracked across cryptography, sensing, and compute strategy and architecture, and leaders will need to focus on building out these systems in order to stay ahead.



Dr. Aaron Kemp
Senior Director, KPMG
Quantum Research Leader



Richard Entrup
Managing Director, KPMG
Head of Emerging Solutions



Michael Isensee
Partner, KPMG
US Leader Cybersecurity
and Technology Risk



For more information, please visit the [KPMG Quantum Page](#).

The shift

Quantum is entering enterprise planning as a transition requirement, not a distant research milestone. Across computing, cryptography, sensing, and converged architectures, regulatory mandates, sovereign capital commitments, workforce formation, and integration decisions are advancing alongside increasingly capable hardware. What is compressing is not the science, but the preparedness window.

Post-Quantum Cryptography (PQC)

What is post-quantum cryptography, in plain English?

Post-quantum cryptography is the effort to replace or upgrade today's encryption before quantum computers can break some of the methods widely used now. In simple terms, it is a security modernization project. It matters today because sensitive data stolen now can be decrypted later, once quantum capability matures — so the stronger protections need to be in place now. And because encryption is buried inside networks, applications, devices, and third-party systems, switching it out takes years.^{8,9}

A real-world example shows why this is already moving.

The U.S. government has already told federal agencies to identify where vulnerable encryption is used, prioritize the highest-risk systems, and estimate what it will cost to migrate.¹¹ CISA has warned organizations to start now because of the “harvest now, decrypt later” risk — the reality that attackers collect encrypted data today and unlock it once quantum capabilities mature.⁹ In May 2026, a draft executive order under review proposed firm deadlines for federal migration of digital signatures and key quantum-vulnerable encryption, signaling that voluntary guidance is moving toward enforceable timelines.¹⁰ Industry analysis published the same month found that the cost of post-quantum migration rises significantly the longer organizations wait, with later movers facing compressed timelines and reduced vendor flexibility.¹¹ This is a current planning issue for any organization with long-lived sensitive data, regulated systems, or complex vendor environments.

Quantum Computing

What is quantum computing, in plain English?

Quantum computing is a different way of solving certain problems that are too complex or time-consuming for today's classical computers to handle efficiently. It is not better at everything. Its strengths are in specific types of work — simulating molecules and materials, exploring many possible options in an optimization problem, or improving parts of complex financial models. That is why no one expects quantum to replace traditional computing. Instead, companies are identifying where quantum adds value alongside classical computing and AI.¹²

A real-world example shows what this looks like.

In May 2026, Q-CTRL, working with IBM Quantum, reported the first evidence of practical quantum advantage on a commercially relevant problem. A materials-science simulation that takes more than 100 hours on industry-standard classical software ran in approximately two minutes on the quantum system — a roughly 3,000-fold speedup, with accuracy matching existing tools.¹³

Days earlier, Cleveland Clinic, RIKEN, and IBM announced they had used quantum-centric supercomputing to simulate a 12,635-atom protein, the largest known biological molecule ever simulated on quantum hardware and roughly 40 times larger than what the same method could handle six months prior.¹⁴ The work has direct implications for drug discovery, where existing computational methods struggle as molecule size grows.



We are at an inflection point where rapid quantum computing advancements, fueled by governments, VCs, and startups, are accelerating the arrival of the quantum era. This progress poses a profound risk to current encryption standards. Consequently, navigating the transition to Post-Quantum Cryptography (PQC) is no longer a distant concern; **it is an urgent business imperative that organizations must address today.**

Richard Entrup *Managing Director, KPMG Head of Emerging Solutions*

Other commercial work is still exploratory, and that is exactly where most enterprises will begin. In 2025, HSBC reported exploratory work with IBM using quantum-enhanced methods for algorithmic bond trading, with the two companies saying the approach improved some trade-fill predictions versus common classical techniques.¹⁵ IBM has continued expanding cloud-accessible quantum systems and hybrid workflows so users can test quantum computing alongside traditional systems rather than in isolation.¹² And the U.S. Department of Energy has funded Quantum Testbed Pathfinder projects to help researchers establish where today's quantum systems are useful in practice.¹⁶

The key point for business leaders is that quantum computing has crossed an important threshold. It is no longer purely theoretical, and in narrow but real-world domains it is now outperforming the best classical alternatives available today.^{13,14} That does not mean buying a quantum computer. It means identifying where hybrid experiments can create insight, speed, or advantage in a few targeted areas.¹⁶

Quantum Sensing

What is quantum sensing, in plain English?

Quantum sensing uses extremely sensitive devices to detect very small changes in motion, location, time, gravity, or magnetic fields. Put simply, it is about measuring the physical world with much greater precision than conventional tools allow. This matters because better measurement leads to better navigation, better diagnostics, and better monitoring of complex environments.

A real-world example shows why this is already moving.

One practical example is navigation when GPS is weak, unavailable, or being jammed. Boeing said it completed a flight test in which an airplane used quantum sensors to navigate without GPS for four hours. The idea is straightforward: instead of relying on outside signals, the system measures the aircraft's own motion with much greater precision. That is why governments and defense organizations are testing quantum sensors now, and why business leaders should watch transportation, logistics, aerospace, and critical infrastructure.²⁸

The key point for business leaders is that quantum sensing is reaching practical use cases sooner than large-scale quantum computing. One of the first business impacts of quantum will come from better measurement and navigation tools, not just faster computers.¹⁷

Hybrid Architecture

What is hybrid architecture, in plain English?

Hybrid architecture means different kinds of computing systems working together instead of trying to make one machine do everything. In this case, it means traditional computing, AI accelerators, and quantum systems each handling the parts of a problem they are best suited for. That matters because quantum is not expected to replace classical computing. As it becomes useful, it is being added into a broader technology stack alongside existing systems.¹⁸

A real-world example shows why this is already moving.

In 2026, IBM published a reference architecture for what it calls quantum-centric supercomputing. The idea is to connect quantum processors with CPUs, GPUs, storage, and orchestration software so they work together in one environment rather than as isolated tools.¹² The Cleveland Clinic and RIKEN protein simulation described earlier is this architecture in practice: IBM quantum processors at two institutions working in tandem with two of the world's most powerful classical supercomputers (Fugaku in Japan and Miyabi-G at the University of Tokyo) to solve a single problem.¹⁴ Government roadmaps point the same way: the U.S. Department of Energy describes computing, sensing, and security as linked pathways, and DARPA's benchmarking initiative is validating the path to utility-scale quantum computing.^{18,19}

The key point for business leaders is that the question is not whether to buy a quantum computer. It is whether your technology environment is flexible enough to plug quantum into existing systems as its useful domains expand. Architecture readiness matters before quantum performance does.

Quantum requires more than computers

Quantum will not become useful at scale because a better machine suddenly appears. It depends on a broader foundation around the hardware: trained people, research centers, software tools, cloud access, secure networks, testing environments, and vendors that can help organizations experiment and integrate these systems into their existing technology environments.

That broader build-out is already underway. In November 2025, the U.S. Department of Energy announced a \$625 million renewal for five National Quantum Information Science Research Centers, explicitly backing not just research but also workforce opportunities, community resources, and industry partnerships across the quantum ecosystem.²⁰ The National Quantum Initiative frames quantum as a national effort spanning science, workforce, industry, infrastructure, economic security, and international cooperation.²¹ The 2026 quantum foundry announcement and the PsiQuantum letter of intent represent the next wave of this build-out, moving from research investment toward domestic manufacturing capacity.^{22,3}

For enterprises, the implication is straightforward. Quantum readiness is not just about whether or when to use a quantum computer. It is about whether the organization has the people, architecture, partnerships, and security foundations to respond as the technology matures. That is why many companies are starting by reskilling teams, testing access through cloud and hybrid environments, and strengthening vendor and security planning before making any major technology bets.

Competitive advantage will come less from owning quantum hardware and more from being ready to use the surrounding infrastructure when it matters.

The cost of inaction

The risk of waiting is not that competitors will suddenly leap ahead with fully mature quantum systems. The bigger risk is that organizations lose the time to prepare for changes that are already showing up in security, vendor requirements, architecture decisions, and workforce planning. Delay does not eliminate the work. It compresses it.

- **Security exposure grows over time.** Data with a long shelf life can be collected now and targeted later, while migration to stronger cryptography takes years across applications, devices, vendors, and networks.^{9,8} Recent breakthroughs in quantum error correction are compressing the timeline further: in mid-2026, QuEra announced a significant improvement in physical-to-logical qubit ratios, a development that brings fault-tolerant quantum computing closer than many migration plans assumed.²³
- **Regulatory and customer pressure can turn a long transition into a rushed one.** Once standards, mandates, or procurement requirements tighten, organizations that delayed inventory and planning will face shorter timelines, higher remediation costs, and more operational disruption.⁸ The May 2026 draft executive order proposing federal PQC deadlines is an early signal that voluntary guidance is becoming enforceable.¹⁰
- **Vendor and architecture choices become harder to unwind.** Companies that do not plan early risk locking in products, contracts, and cloud environments that are slower or more expensive to adapt later, especially as quantum-readiness expectations begin to appear in refresh cycles and supplier requirements.²⁴
- **Capability gaps widen.** Early movers are building experience through pilots, partnerships, and reskilling. Organizations that wait will find it harder to access talent, shape vendor roadmaps, or influence how emerging standards are applied in their industry.²⁰
- **Strategic flexibility erodes.** The longer leaders wait, the fewer low-cost options remain. What starts as a manageable preparedness effort can later become a forced response driven by risk, regulation, or market pressure.

The goal is not to make large bets too early. It is to build enough understanding, inventory, and flexibility now so the organization can move deliberately later, instead of reacting under pressure.

A Unified Roadmap: Actions for Now, Near, and Next

Elevate governance

- Bring quantum into executive discussions across security, technology, procurement, and risk.
- Treat encryption timelines, vendor risk, geopolitics, and talent as business decisions, not just technical ones.

Map strategic exposure

- Identify where long-lived data, third-party dependencies, IP, and regulated systems create risk.
- Factor in regulation, ecosystem fragmentation, and talent shortages when planning ahead.

Sequence capital deliberately

- Align quantum investments with existing modernization efforts.
- Move in stages. Build capability without locking into early bets or rushed timelines.

Design for substitution

- Build systems that can adapt, especially around encryption and vendors.
- Keep flexibility as technologies and standards continue to evolve.

Quantum is moving from a future concern to a present-day priority. Organizations that stay flexible, build internal capability, and plan early will be in a stronger position when change accelerates.

How leading companies are already responding

As institutions move faster — and as the technology matures alongside them — leading companies are treating quantum as a business and risk issue, not just a research topic.

This shows up in a few clear ways:

- Post-quantum cryptography is being built into modernization plans now, not pushed off as a future compliance task.⁸
- Small hybrid pilots, often tied to AI and high-performance computing, are demonstrating how quantum and classical systems work together before standards are fully defined.^{15,25}
- Organizations are tracking quantum sensing, especially where it stands to improve navigation or diagnostics.^{26,18}
- Teams are investing in skills through reskilling and partnerships with universities.²⁰
- Vendor decisions increasingly factor in geopolitics and supply-chain risk, not just performance.
- Boards are beginning to include quantum in risk discussions and long-term planning.²⁷

Taken together, this points to a clear shift. Quantum is now shaping decisions around security, systems, vendors, and long-term investment.



Quantum is actionable now. This isn't science fiction.
You've got to start experimenting!

Dr. Aaron Kemp, Senior Director, KPMG, Quantum Research Leader

Scenarios and signals

Quantum development is not linear. Breakthroughs, delays, and geopolitical shifts will shape timing, scale, and enterprise exposure. The notable change in 2026 is that several scenarios once treated as hypothetical are now visible in the evidence.

Scenario	What it means for you
Fault-tolerance arrives faster than planned. Advances in qubit coherence and error correction are compressing hardware timelines. QuEra’s error-correction milestone and Q-CTRL’s practical quantum advantage result, both in mid-2026, show this is already underway. ^{13,23}	Hardware acceleration narrows the preparedness window.
Hardware progress slows. Commercial systems would stay confined to specialized domains for longer — yet regulatory mandates and capital mobilization advance regardless of hardware maturity.	The preparedness work pays off in either case.
Migration timelines extend. Enforcement will vary across sectors and geographies.	Data-longevity risk persists under harvest-now, decrypt-later dynamics, so the case for early inventory does not weaken.
Sensing outpaces computing. Field deployments precede broad computational advantage, as the Boeing flight test suggests.	The first competitive shifts may arrive through sensing.
Convergence reshapes architecture first. Hybrid models are already influencing system design, with IBM advancing integrated High-Performance Computing (HPC), and quantum data-center architectures. ¹²	Architectural literacy comes before raw performance.



Our goal at IBM is to bring useful quantum computing to the world. To do that, we are working with partners from across industry to push the boundaries of quantum-centric supercomputing — **where quantum processors work in concert with classical high-performance computing to solve problems that were previously out of reach.** This hybrid architecture is the blueprint for achieving applications with a quantum advantage now, and reaching large-scale, fault tolerance before the end of the decade.

Dr. Scott Crowder, Vice President, IBM Quantum Adoption

Conclusion

Quantum has passed a threshold. It is not yet pervasive the way cloud, mobile, and AI became pervasive, but it is now delivering specific, commercially relevant results today — and the frontier is expanding faster than many migration plans assumed.

Institutional momentum built ahead of technological certainty. National strategies, capital commitments, and emerging cryptographic mandates were already shaping how organizations think about risk and infrastructure. In 2026, for the first time, the technology is catching up. Practical quantum advantage has been demonstrated on a commercially relevant problem. Biologically meaningful molecules are being simulated at scales unthinkable a year ago. Error correction is advancing faster than many migration plans assumed. The window between preparation and consequence is compressing.


The practical work starts now:

- understand where your encryption is exposed;
- build systems that can adapt;
- ask more of your vendors;
- run targeted pilots to build familiarity;
- monitor sensing where it intersects with your business;
- begin training the teams that will manage the transition.

The companies that benefit will not be those that move first on large-scale quantum bets. They will be the ones that prepare early and preserve flexibility as the landscape evolves.

FRONTIER JOURNEY
Building your Blueprint

Download the **Quantum Frontier Blueprint**
to help navigate your journey.



NextGen Computing

Quantum is one specialized architecture among many. The same forces driving enterprises to think about quantum readiness — workloads outgrowing general-purpose hardware, energy and memory becoming hard constraints, and standards emerging that let different systems interoperate — are reshaping the broader computing stack. Quantum gets the headlines, but the more immediate shift for most enterprises will come from neuromorphic, photonic, processing-in-memory, and other architectures advancing in parallel and converging into a single, more heterogeneous computing portfolio.

The recent public debut of Cerebras is a reminder that next-generation computing is no longer a research story alone.¹ Investors are beginning to treat specialized computing architectures as a real strategic category, not a technical curiosity. That signals a broader shift already underway: enterprise computing is moving away from one-size-fits-all systems and toward a portfolio of architectures optimized for different kinds of workloads.

For decades, enterprise computing rested on two assumptions: that computing power would keep improving steadily, and that companies could rely on simple, all-in-one systems to run their operations. Both are now breaking down. Moore's Law is delivering less predictable economic benefit, while newer workloads demand architectures designed for speed, efficiency, or adaptability rather than general-purpose scale alone.

As computing becomes more diverse, traditional all-in-one systems can limit flexibility rather than simplify it. Companies are moving toward a mix of systems that need to work together, and the ability to coordinate them is becoming a key advantage.

Why this matters now

The pressures are already showing up in enterprise decisions. AI inference is becoming a larger and more persistent cost center, memory and power constraints are harder to ignore, and organizations are finding that not every workload benefits from the same architecture. The strategic question is no longer whether a single platform will dominate, but how to place workloads across a more diverse computing environment without creating unnecessary complexity.

AI as the Frontier Multiplier

AI workloads are pushing traditional computers to their limits. Next-generation architectures — neuromorphic chips, photonic systems, and in-memory computing — each solve one specific problem better than CPUs or GPUs can. The multiplier effect happens when organizations combine these specialized systems with traditional computing, routing different parts of the AI task to whichever system is most efficient.

The shift

Three forces are reshaping enterprise computing.

- **Hardware alone is no longer enough.**

For decades, chip improvements drove computing gains directly. Today, hardware advances alone deliver less predictable economic returns. Research from MIT scholars Brynjolfsson, Rock, and Syverson shows that computing technologies require complementary investments—in software, talent, and organizational change—to realize economic value. Waiting for faster chips is no longer enough.¹⁶

- **Workloads are outgrowing general-purpose systems.**

Modern workloads — AI, simulations, optimization, and real-time data processing — are fundamentally different from the tasks traditional systems were designed for. General-purpose processors can still run them, but often inefficiently and at higher cost. Specialized hardware is already common in many environments today, particularly in AI.

- **Infrastructure is unbundling**

In the past, systems were delivered as tightly integrated packages combining compute, storage, and networking. That made them easy to buy and deploy, but created long-term lock-in. Today, infrastructure is becoming more modular. Systems can be broken into parts and used more flexibly, especially in high-performance environments, a trend enabled by emerging interconnect standards like Compute Express Link (CXL) and increasingly productized by vendors.

In practice, this lets organizations:

- change hardware vendors without rebuilding everything;
- run different types of computing side by side;
- introduce new technologies over time instead of all at once.

This shift is critical. Without more flexible infrastructure, different types of computing cannot operate together at scale.

The computing architecture landscape

Next-generation computing is not one technology. It is a set of approaches; each designed for specific types of problems and advancing at different speeds.

- **Neuromorphic Computing:** Unlike traditional step-by-step processors, neuromorphic computing mimics the human brain's neural networks to spot patterns faster with minimal power. A prominent example is Intel's Loihi 2 research chip, which packs up to one million artificial neurons. By only consuming electricity when processing an active event, it runs tasks at a fraction of the energy of standard processors, making it hundreds of times more energy-efficient.²

WHAT THIS UNLOCKS: This unlocks true learning at the edge with extreme energy efficiency. Today's edge AI mostly runs pre-trained models. Neuromorphic chips could allow a factory robot or autonomous drone to learn and adapt to its environment in real-time using a fraction of the power, without needing to send data to the cloud. This creates an opportunity space for truly independent, long-duration autonomous systems in power-constrained environments where continuous learning is critical.

MATURITY: Early commercial chips exist; adoption is still concentrated in pilots and research environments.

- **Photonic Computing:** Systems that use light instead of electricity to move and process data, improving speed and reducing energy use in certain operations. Lightmatter, the most visible commercial player, announced its Passage L200 and M1000 co-packaged optics platforms for AI data centers and a record 1.6 Tbps per fiber demonstrated with Qualcomm. The technology directly addresses the interconnect bottlenecks that limit AI infrastructure scaling.³

WHAT THIS UNLOCKS: This unlocks a path to overcoming the data movement bottlenecks and power constraints that are strangling AI scalability. As AI models grow, the energy spent simply moving data between memory and processors is becoming the primary limiting factor. By using light instead of electricity, photonic systems could enable the creation of massive AI models that are currently economically and physically unfeasible. The business opportunity is not just faster AI, but a sustainable path to the next order-of-magnitude increase in AI capability.

MATURITY: Transitioning from research into early commercial development and targeted deployment.

- **Processing-in-Memory (PIM):** Systems that perform computation inside or beside the memory itself, reducing the need to move data back and forth between memory and processor. Major memory manufacturers have begun productizing PIM-enabled high-bandwidth memory (HBM) for AI inference workloads, with early deployments appearing in hyperscaler and research environments.

A related, more advanced variant — analog in-memory computing (AIMC): Performs the AI math directly inside the memory array in the analog domain rather than digitally, which can sharply cut the energy used for AI inference. IBM offers the clearest commercial signal here, with a 64-core mixed-signal analog AI chip (its HERMES project) and 2025 research extending analog in-memory designs to larger models; the main open question is which underlying memory material — such as Resistive RAM (RRAM) or phase-change memory (PCM) — ultimately wins out.⁴

WHAT THIS UNLOCKS: This unlocks a radical reduction in the cost and latency of AI inference, which currently accounts for the vast majority of operational AI workloads. By performing calculations directly inside the memory, PIM and AIMC architectures could make sophisticated AI cheap enough to embed in every device, from consumer electronics to individual industrial sensors. This creates a business opportunity for pervasive intelligence, shifting AI from a centralized, high-cost resource to a ubiquitous, low-cost utility.

MATURITY: Digital PIM is in early commercial experimentation; analog in-memory has working chips but remains largely pre-deployment.

- **Thermodynamic Computing:** Chips that harness the natural electrical noise in their circuits — rather than suppressing it — to perform the random sampling that underpins many AI and optimization tasks, at far lower energy than conventional processors. The clearest commercial signal is the startup Extropic, which published its technical approach and announced its first hardware platform, the XTR-0, in late 2025, with a production-scale chip (the Z1) targeted for 2026. Its Thermodynamic Sampling Units (TSUs) use probabilistic bits (“p-bits”) to draw samples directly from energy-based models. The headline claim is up to roughly 10,000 times better energy efficiency than GPUs for these tasks — a figure that still needs to be proven at production scale.^{5,6}

WHAT THIS UNLOCKS: This unlocks leaps in energy efficiency for an advanced AI workloads at massive scale, creating an opportunity space for applications that are currently unthinkable due to their prohibitive energy requirements, such as large-scale material science discovery or real-time financial risk simulation.

MATURITY: Prototype hardware shipping to select partners; broad commercial use still ahead.

- **Probabilistic Computing:** The broader family that thermodynamic computing belongs to: approaches that trade perfect precision for speed and energy savings where exact accuracy is not required. Thermodynamic chips are its most commercially concrete example today, but the wider category also spans approximate-computing techniques whose activity remains concentrated in venture-backed startups and academic spinouts, with broader enterprise application still emerging.

WHAT THIS UNLOCKS: This unlocks the ability for machines to make fast, effective, “good enough” decisions under real-world uncertainty, much like humans do. Instead of requiring perfect data and delivering a single, precise answer, these systems provide rapid, energy-efficient judgments in messy, unpredictable environments. This creates a business opportunity for more robust and resilient edge AI systems, such as autonomous vehicles navigating in poor weather or medical devices making preliminary diagnoses from imperfect sensor data. It represents a shift from brittle, deterministic AI to more adaptable, real-world intelligence.

MATURITY: Mostly research and early experimentation, apart from the thermodynamic-hardware efforts noted above.

- **Also on the Horizon:** Three further approaches are worth tracking but remain earlier and longer-dated. Reversible, or “near-zero energy,” computing recycles the energy of each calculation instead of releasing it as heat; the UK startup Vaire Computing demonstrated net energy recovery on its first test chip, “Ice River,” in 2025.⁷ Superconducting (cryogenic) computing runs on materials cooled to near absolute zero, where electrical resistance disappears, promising large efficiency gains — a focus of research groups such as IMEC.⁸ DNA data storage encodes information in synthetic DNA for extremely dense, long-lived archival storage, an effort led by Microsoft with the University of Washington and the industry-backed DNA Data Storage Alliance.^{9,10} None are close to enterprise deployment, but each targets a constraint — energy, heat, or storage density — that today’s computing is straining against.

Taken together, these new computing paradigms signal a fundamental break from the one-size-fits-all model that has defined the last several decades. In the medium term, this means enterprise computing will evolve into a heterogeneous portfolio where workloads are strategically matched to the optimal architecture.


For leaders, this is not a mere technical upgrade; it is a profound shift in what is possible. Architectures like Processing-in-Memory could unlock pervasive, low-cost AI, making it cheap enough to embed in every device and process. Photonics offers a path to sustainable, large-scale AI, overcoming the power and data-transfer bottlenecks that currently limit growth. Neuromorphic and Probabilistic computing will enable truly autonomous systems that can learn and make effective decisions in messy, real-world environments without constant connection to the cloud.

The opportunity is the chance to create entirely new business models built on ubiquitous intelligence, to accelerate scientific discovery in domains currently blocked by computational cost, and to build resilient, next-generation products that are impossible with today’s computing stack.

The central message for leaders is that a change in computing architecture is a direct catalyst for creating entirely new opportunity spaces.

Indicative Maturity Bands

Indicative bands based on current research activity, commercial signals, and enterprise adoption patterns.



Signals under the radar

Next-generation computing is advancing through targeted investment, research programs, and emerging interoperability standards rather than broad enterprise adoption — which means the architectural choices being locked in now will shape what enterprises can build a decade out.

National labs are funding designs built to get past the energy and adaptability limits of today's AI. Researchers at Los Alamos, working under the CHIPS and Science Act of 2022, are developing neuromorphic systems modeled on the human brain, aiming for AI that processes information through feedback loops rather than the rigid, one-direction logic of today's large language models.¹¹ Oak Ridge runs Frontier, the world's first exascale supercomputer, alongside a Quantum Computing User Program that gives researchers access to working quantum systems — signaling where neuromorphic and quantum architectures can scale once they leave the lab.¹²

Major technology providers are placing bets across multiple computing approaches at once. IBM continues to invest in neuromorphic research,¹³ while Nvidia made a \$20 billion deal in December 2025 to license Groq's inference technology and bring its founding team in-house — by far the largest transaction in Nvidia's history.¹⁴ The message across these moves is consistent: no single architecture is going to own the next decade.

The standards layer is quietly preparing for this mix. CXL, in particular, is making it easier for different types of computing chips — CPUs, GPUs, accelerators, and newer designs — to share memory and work together in the same system.¹⁵

Risks and tradeoffs

A more heterogeneous computing model creates real advantages, but it also introduces new forms of complexity. Integration can become harder as organizations manage more vendors, software stacks, and orchestration layers. Tooling and standards are still maturing in several areas, which raises the risk of fragmented pilots that never scale. Talent is another constraint: success depends not only on access to new hardware, but on the ability to map workloads correctly, manage data movement, and govern performance, resilience, and security across mixed environments. The risk is no longer only lock-in to a single platform; it is also lock-in to poorly integrated complexity.

What leaders are doing

Leading organizations are approaching next-generation computing as a portfolio strategy rather than a single platform decision.

- Running controlled pilots**
 Siemens and NVIDIA's first manufacturing blueprint at the Siemens Electronics Factory in Erlangen, Germany is a clear example: the companies are running a fully AI-driven adaptive manufacturing pilot at one site before scaling.¹⁷ Mayo Clinic partnered with Cerebras to build a genomic foundation model combining patient records with genomic data, starting with a focused application in rheumatoid arthritis treatment selection before expanding.¹⁸
- Designing for heterogeneous compute environments**
 AMD released flagship server processors with advanced interconnect support, making it a foundational feature of modern servers. NVIDIA's NVLink Fusion platform later allowed companies to combine NVIDIA chips with custom silicon from partners, explicitly built for heterogeneous systems rather than single-chip stacks.^{19,20}
- Maintaining vendor flexibility**
 In Europe, member states and 19 companies formed a consortium (IPCEI-CIS – 8ra) to develop next-generation, multi-vendor computing use cases collectively, giving members the ability to learn across multiple providers without being captured by any one.²¹
- Building cross-disciplinary capability**
 Bosch established the Bosch Center for Artificial Intelligence (BCAI), a dedicated internal organization that employs hybrid and probabilistic modeling, neuro-symbolic AI, and other experts to partner with Bosch's business units — from manufacturing and supply chain to automotive components — embedding deep expertise directly into industrial domain problems.²²

From insight to action

The strategic question for enterprise leaders is not which single computing architecture will win. It is how to build an operating model that can match workloads to the right architecture, capture performance and cost advantages where they are real, and preserve enough flexibility to adapt as the landscape evolves. That makes next-generation computing less a procurement exercise than a portfolio and orchestration challenge.

In the near term, organizations should focus on three priorities.

1

Map workloads to architectural fit

Identify where specialized systems create a measurable advantage in cost, latency, energy efficiency, or capability, and distinguish those cases from workloads that should remain on general-purpose infrastructure.

2

Build integration capability early

Invest in the internal ability to test, orchestrate, and govern mixed environments so pilots can scale and new architectures can be introduced without fragmenting the broader technology stack.


3

Preserve strategic flexibility

Avoid infrastructure choices that assume one architecture will dominate. The enterprises best positioned for advantage will be those that combine multiple computing approaches effectively while keeping vendor, standards, and deployment options open.

FRONTIER JOURNEY
Building your Blueprint

Download the **NextGen Computing Blueprint** to help navigate your journey.





Future of Trust

Architecting trust in a zero-trust world

We are entering a world in which enterprises can no longer assume that what they see is real or that what a system tells them is true. This zero-trust environment is the new reality, driven by two forces: the rise of highly realistic synthetic media and the increasing autonomy of AI agents.

Consider an AI system that reports, “I received the appropriate approval for this transaction.” In a zero-trust world, that statement alone is not proof. Instead, leaders must demand independent, verifiable evidence that the approval actually occurred. The same is true for digital media, where the inability to distinguish authentic from synthetic content presents a direct risk to decision-making, brand reputation, and fraud prevention.

This introduces the fundamental question for the enterprise: In a world where you potentially cannot trust anything by default, how do you establish enough trust to transact, make decisions, and operate?

The answer cannot be found in traditional governance overlays. Instead, trust must become a core architectural requirement of the enterprise technology stack. This means building systems that provide independent verifiability, cryptographic auditability, and clear data provenance, all without creating a terrible user experience that slows down the pace of business. The challenge is re-architecting the enterprise for a world where trust must be explicitly built, not assumed.

New approaches are emerging to track the origin and lineage of data and content, with standards such as Coalition for Content Provenance and Authenticity (C2PA) enabling verifiable provenance to be attached directly to digital assets.¹ Governance frameworks, including the National Institute of Standards and Technology AI Risk Management Framework, the EU AI Act, and ISO/IEC 42001, are formalizing this shift by emphasizing lifecycle risk management, monitoring, and structured oversight of AI systems.^{2,3,4}

AI as the Frontier Multiplier

How is AI affecting the way people trust your business? Digital credibility is as important as ever as consumers try to navigate what is real. Identity and authentication, transactional trust, and system integrity are all areas where AI can be both the problem and solution. The Future of Trust will require new guardrails for how organizations use the technology.



Matt Johnson
AI Audit and Assurance Leader



Bryan McGowan
Principal, Global Trusted AI Leader



Mike Sawyer
Partner, US Audit Leader for AI & Innovation



Why now?

Four powerful forces are converging to make technical trust a board-level priority:

- 1. The Scale of Automation:** The rise of agentic AI promises to deploy new kinds of artificial actors across workflows at scale, from customer interactions to supply chain.
- 2. The Rise of Synthetic Reality:** Generative AI has expanded access to highly realistic synthetic content, with documented enterprise impact – including deepfake fraud attempts in contact centers that rose by more than 1,300% in 2024, jumping from an average of one per month to seven per day, with contact centers facing projected fraud exposure of \$44.5 billion in 2025.⁵ The Stanford AI Index continues to document rapid improvements in model capability and accessibility.⁶
- 3. The Distributed Ecosystem:** More and more, enterprises must operate across multi-cloud environments, partner ecosystems, and global supply chains, where workflows depend on shared infrastructure and third-party systems.
- 4. The Convergence of Regulation and Risk:** New AI governance frameworks and cybersecurity requirements are increasingly requiring organizations to demonstrate how systems are governed and evaluated. This shifts assurance from a best practice to a requirement, reinforcing the role of external validation even as system-level verification capabilities advance.^{3,4}

These dynamics are not entirely new. When IoT emerged, enterprises faced similar architectural questions around identity, data provenance, and distributed verification. The governance frameworks, technical standards, and trust models developed for connected devices are now being extended for agentic AI, with the critical difference that AI agents make autonomous decisions and generate synthetic content, multiplying the verification burden.

Why it matters

Trust failures shift from isolated incidents to systemic risks that can disrupt entire value chains. Operating a modern digital enterprise depends on verifiable trust.

- Operational Risk Requires Verifiable Settlement**
In financial services, institutions must ensure that tokenized assets and programmable transactions operate within trusted settlement systems to prevent market-destabilizing errors or fraud. BlackRock's BUIDL fund has grown significantly across multiple blockchains since its March 2024 launch, illustrating the scale at which trusted settlement infrastructure now operates.⁷
- Supply Chain Traceability for Safety**
In healthcare and life sciences, organizations must maintain traceable data and biological supply chains to ensure patient safety and regulatory compliance.
- Autonomous Safety Accountability**
In industrial sectors, companies deploying autonomous machines must be able to demonstrate safety, reliability, and operational accountability to avoid catastrophic physical and financial damage.

The core requirement is reliable coordination across complex networks of partners, platforms, and infrastructure.

Organizations that can verify identity, data, and system behavior gain three advantages: reduced systemic risk, lower coordination friction across ecosystems, and the ability to operate new trust-dependent business models.

The shift: from institutional assurance to technical verification

Verification is moving from periodic checks to continuous validation integrated into system operations. As more activity is executed by software, platforms, and autonomous systems, verification must occur within the flow of operations rather than through periodic checks.

This transition is reflected in four structural shifts:

- **Identity as the primary security boundary**
Access and activity are verified through identity rather than network location. Identity now extends beyond employees to include devices, software services, and AI agents.
- **Provenance of data and content**
Systems are tracking the origin and lineage of data, media, and digital assets to determine whether information can be trusted.
- **Continuous assurance of systems**
AI models and automated systems are monitored and evaluated on an ongoing basis to ensure reliability, performance, and accountability.
- **Verification in transactions**
Transactions are validated at the point of execution through cryptographic methods, reducing reliance on reconciliation and post-processing.

As trust becomes part of core system architecture, control over how it is defined and enforced becomes a source of strategic advantage. Organizations that shape identity, transactions, and system interactions will influence how value and risk are distributed.

In practice, the future will require a mix of both embedded controls, internal validation systems, and external assurance. External validation will remain an important component of trust, especially in regulated environments and across more and more complex environments. At the same time, as systems operate at greater speed and scale, more verification will need to happen within the system itself.

A shift in mindset: From governance overlay to architectural requirement

For years, trust has been treated as a governance function, a compliance checklist applied after a system is built. That model is shifting. In a zero-trust world, trust must be a core architectural requirement, engineered into systems from the start through verifiable identity, data provenance, and cryptographic assurance.

The three domains of technical trust

To manage this new landscape, leaders should focus on building capabilities across three interconnected domains.

1. Identity and Authenticity (Who and What is Real?)

Capabilities that establish who or what is participating in a digital interaction and whether information can be trusted. As digital activity scales and synthetic content becomes more prevalent, organizations are shifting from assumed identity to verifiable identity, and from accepted content to provable origin.

Organizations are implementing verifiable credentials and decentralized identity models to enable secure cross-organizational identity federation, with 62.3% of surveyed North American enterprises allocating budget for verifiable credentials evaluations in 2025-2026.⁸ Enterprise platforms such as Microsoft's Entra Verified ID illustrate this shift, though adoption is uneven across industries. Content provenance standards such as C2PA are establishing ways to attach verifiable origin data to digital media as AI-generated content becomes more prevalent.¹

2. Cryptographic and Transactional Trust (Verifying Value in Motion)

Capabilities that secure and validate transactions and digital assets, enabling trust to be established within the transaction itself rather than relying on post-transaction reconciliation or intermediaries.

Adoption is most visible in financial services. Asset managers such as BlackRock are launching tokenized funds on similar infrastructure.⁹ Beyond finance, techniques such as zero-knowledge proofs are being explored to allow verification of data or transactions without exposing underlying sensitive information.

3. System Integrity and Accountability (Ensuring Systems Behave as Intended)

Capabilities that ensure automated systems operate reliably and can be governed at scale. Formal AI assurance practices are taking shape, including structured model evaluation and adversarial testing to identify failure modes and vulnerabilities. Organizations are preparing for longer-term infrastructure risks, including the transition to post-quantum cryptography. In August 2024, NIST released three finalized post-quantum encryption standards designed to withstand future quantum attacks.¹⁰

Implications for leaders

Designing trust into system architecture changes how organizations build and operate. Identity, verification, and accountability are becoming part of core infrastructure rather than standalone controls.

Trust can no longer sit within a single function. It requires coordination across technology, security, and data leadership, supported by systemic, internal, and external approaches.

Leaders should focus on two areas:

1. Readiness for scaled automation

Can the organization verify identity, data, and behavior at a level required to support its AI and automation plans?

2. Exposure to trust-related failure

Does the organization understand the financial and reputational impact of failures in identity, data integrity, or system accountability, and are mitigation measures in place?

Trust adoption follows adoption curves, not exponential trajectories. Society already delegates critical decisions to complex systems like banking, aviation, and power grids – despite known failure modes. When failures occur, trust dips but doesn't disappear. Organizations building trust infrastructure today are positioning themselves for the inevitable recovery and scale that follows. This is a marathon, not a sprint. Organizations that build verification into their systems can scale with confidence.

A Unified Roadmap: Actions for Now, Near, and Next

Leaders should operationalize trust as a core component of system design. The path forward is not a single initiative, but a staged transition from fragmented controls to integrated verification across systems, transactions, and ecosystems.

Now

Establish an enterprise identity layer

Extend identity across employees, systems, devices, and AI agents, starting with high-risk workflows and key partner integrations.

Implement provenance and AI assurance in priority areas

Track data lineage and content origin in high-risk workflows. Stand up model evaluation, monitoring, and audit processes.

Assess trust dependencies and cryptographic exposure

Identify where identity, verification, and encryption are critical, and where future risks (including post-quantum) are highest.

DECISION: Where does the organization need verifiable trust first—and which processes create the greatest risk if identity, data, or system behavior cannot be validated?

Near

Embed verification into core workflows and transactions

Validate identity, data, and system actions at the point of execution (e.g., customer transactions, partner data exchanges, automated decisions).

Operate continuous assurance across systems

Monitor and test models, data, and workflows in real time, with clear ownership and escalation when issues are detected.

Extend trust to key partners and platforms

Apply consistent identity, provenance, and verification standards across vendors, ecosystems, and third-party systems.

DECISION: Which elements of trust should be owned internally versus integrated from external platforms—and how will that shape control over data, transactions, and system behavior?

Next

Build trust into core infrastructure

Make identity, verification, and assurance standard services used across all systems, not separate controls applied after the fact.

Scale operations that rely on built-in verification


Expand autonomous systems, tokenized assets, and shared data environments where validation happens automatically within workflows.


Adopt next-generation cryptography

Implement post-quantum encryption and privacy-preserving techniques to secure long-term data and transactions.

FRONTIER JOURNEY
Building your Blueprint

Download the **Future of Trust Blueprint** to help navigate your journey.





Space

If recent headlines about orbital compute sound exaggerated, that may be the point: the technology may still be emerging, but the strategic implications are arriving faster than most leadership agendas. What is changing is not just the economics of launch or the novelty of new space ventures; it is the growing role of space as a foundational layer of enterprise infrastructure—powering connectivity, timing, sensing, intelligence, and resilience across industries. Whether every near-term promise materializes or not, the signal is unmistakable: space is shifting from a niche sector story to a board-level question about competitive advantage, operational dependence, and long-term control over critical infrastructure.

The space economy is at a strategic inflection point. Historically, the domain of governments and specialized aerospace firms, space is now becoming a foundational layer of the global economy.

Space-based capabilities support navigation, communications, supply chains, financial transactions, and environmental monitoring, embedding space directly into core business operations. As a result, most organizations are already dependent on space infrastructure, often in ways that are indirect and not well understood.

For enterprise leaders, the question is not whether they are in the space economy, but where it intersects with their business. This often begins with understanding reliance on space-enabled services and identifying where space-derived data can improve operations or create new value. For boards and C-suites, the priority is to manage this dependence as a strategic issue. This includes governing critical dependencies, mitigating emerging risks, and identifying where space-based capabilities can create advantage.

The broader market for space-related goods and services is projected to grow from approximately \$630 billion in 2023 to more than \$1.8 trillion by 2035.¹

What Leaders Should Be Looking for in the Space Economy

- **Revenue Growth:** Creating new services using satellite and geospatial data, such as climate-risk analytics for insurers, precision agriculture insights, or supply chain visibility products.
- **Cost Avoidance:** Using satellite intelligence to identify shipping delays, monitor infrastructure remotely, reduce downtime, and improve fuel or resource efficiency.
- **Risk Exposure:** Assessing how disruptions to GPS, satellite communications, or Earth observation data could affect logistics, financial systems, operations, or customer service.
- **Competitive Differentiation:** Embedding persistent global sensing and connectivity into products, operations, and AI systems faster than competitors.

AI as the Frontier Multiplier

AI is accelerating the value of space in two directions at once. On the ground, AI models are turning raw satellite data into real-time operational intelligence, compressing what once took days of analysis into decisions made in seconds. In orbit, the scale of AI's compute and energy demands on Earth is pushing companies to explore processing data where it is generated. The result is a feedback loop: AI makes space data more valuable, and space is becoming part of the infrastructure AI runs on.



Brian Miske
Principal, KPMG
Americas Space Leader



The new value stack: translating space into profit and loss impact

The challenge is not awareness of space, but a clear line of sight to its impact on enterprise performance. To bridge this gap, organizations can view space through a new Value Stack (Exhibit 1). This framework moves beyond abstract market sizing to quantify the opportunity and risk in terms executives and board can act upon.

EXHIBIT 1

Value Stack Component	Description & Enterprise Impact	Enterprise Impact
Revenue Growth	Create new products or premium services using satellite data. For example, a company could use Earth imagery to help insurers price risk more accurately, help farmers monitor crop conditions, or help logistics teams track disruptions and reroute shipments faster. ²	New products and data services
Margin Expansion	Cut costs by making better day-to-day decisions. Satellite data can help companies spot supply chain delays earlier, schedule maintenance before equipment breaks, and use fuel, water, or other resources more efficiently. ³	Operational efficiency
Capital Efficiency	Commercial providers increasingly offer space-enabled capabilities as subscription-based services and APIs. This allows enterprises to adopt advanced capabilities without building proprietary infrastructure or specialized aerospace expertise. ⁴	Faster deployment and lower complexity
Risk Mitigation	Reduce the business impact if space-based services go down. If GPS, satellite communications, or Earth observation data are disrupted, companies can face shipping delays, payment issues, service outages, or poor decisions. Planning backups and alternative providers helps protect revenue and customer trust. ⁵	Business continuity and resilience

Viewing space through the lens of this type of economic framing helps it become a measurable driver of financial performance.




Space is no longer a distant frontier - it is becoming a strategic layer of enterprise infrastructure, and leaders need a clear point of view on where it creates advantage, dependence, and risk.

Brian Miske, Principal, KPMG, Americas Space Leader

A framework for execution: operating model and governance

Organizations may agree that they should be focusing on space, but the question is often, “Who inside my company owns this?” A lack of defined ownership and governance is one of the barriers to action.

Leading enterprises are not creating large, new “space departments.” Instead, they are embedding space as a horizontal capability layer managed within existing business functions (Exhibit 2).



While the operating model defines how the work gets done, formal corporate governance is required to elevate space from an optional innovation to a core business priority. Because space is a critical infrastructure, its governance should be integrated into the highest levels of corporate decision-making. A best-practice governance model can include:

- **Accountability: A Single C-Suite Owner:** Final accountability for the enterprise space strategy must reside with a single C-suite executive, typically the Chief Technology Officer or Chief Strategy Officer. This ensures clear ownership, prevents diffusion of responsibility, and provides a single voice to the investment committee and the board.
- **Oversight: The Enterprise Risk Committee:** Dependency on space infrastructure is a significant and growing business risk. As such, oversight should belong to the board's Enterprise Risk Committee, placing space-related risks on par with other top-tier risks like cybersecurity and supply chain failure.
- **Funding: The Technology Investment Committee:** Space initiatives should not be funded out of a discretionary innovation budget. Instead, funding requests should be routed through the standard Technology Investment Committee, where they can be evaluated with the same rigor as any other major investment in core enterprise infrastructure like cloud platforms or ERP systems.

This combined framework of an operating model and a governance model ensures that space strategy is not only tied to business outcomes but is also embedded into the core accountability, risk, and capital allocation processes of the enterprise.

Reference architecture: integrating space into the enterprise stack

For Chief Information and Technology Officers, the strategic “why” must be followed by the technical “how.” Before funding can be allocated, leaders need to visualize how space-based capabilities integrate into the existing enterprise technology stack. This is not about building a new, isolated system, but about treating space as a powerful new data source for the infrastructure you already run.

The right image reference architecture illustrates this data flow from orbit to decision.

As the diagram shows, data flows from the Orbital Layer (satellites), through a commercial Platform Layer (partner APIs), and directly into your existing Enterprise Layer. The key takeaway for a CTO is that space-derived data becomes another trusted input for your core systems—cloud, edge, ERP, and AI models—augmenting their capabilities without requiring a fundamental redesign of your architecture.



THE HARDWARE

Layer 1 — Space Access and Infrastructure

This is the physical foundation: launch systems, satellites, and orbital platforms; today's stations; and tomorrow's commercial outposts. It is the layer that makes the rest possible by putting capability where it can operate in orbit and beyond.

THE NETWORK

Layer 2 — Space Data and Connectivity

Once the hardware exists, it becomes an information utility: communications, Earth observation, and positioning/navigation/timing signals that deliver connectivity, visibility, and synchronization across the global economy. This is where space stops being a destination and becomes a service.

THE BRAIN

Layer 3 — Enabling Technologies

This is the intelligence layer: artificial intelligence (AI), edge computing, quantum sensing, advanced materials, autonomy, and the architectures that make the stack faster, more resilient, and more scalable. Layer 3 is where space systems become increasingly software-defined, capable of adapting prioritizing, and acting with less human intervention.

THE APP STORE

Layer 4 — Space Access and Infrastructure

This is where value becomes tangible. It's the farmer optimizing irrigation, the logistics operator tracking movement in near-real time, the bank depending on precise timing, the insurer validating catastrophe exposure, and public agencies strengthening disaster response. Layer 4 is where the space stack becomes “just how the world works.”

Signals Under the Radar: AI Is Pulling Space Into the Infrastructure Conversation

Growing AI demand is putting pressure on energy systems, land, cooling, and data center capacity on Earth. As a result, companies are beginning to explore new ways to connect computing, communications, and infrastructure across both ground and orbital systems.

Several emerging dynamics are worth monitoring:

- AI growth is increasing demand for power, connectivity, and compute infrastructure. Telecom, cloud, and space providers are starting to work together on hybrid ground-space networks and services.
- Companies are becoming more focused on resilience and flexibility, avoiding dependence on a single infrastructure or connectivity provider.

Most enterprises will not move workloads into space anytime soon. But the broader signal is important: space is increasingly becoming part of the long-term infrastructure stack supporting connectivity, compute, resilience, and operational continuity.

A Unified Roadmap: Actions and Technologies for Now, Near, and Next

An effective strategy requires not only sequencing actions but also aligning them with the maturity of the underlying technologies. This unified roadmap provides a clear framework for what to do and which technologies to use over a multi-year horizon.

STRATEGIC ACTIONS & KEY DECISIONS

Now

MAP DEPENDENCIES & RUN PILOTS

- Conduct a "space dependency audit" to identify where critical operations already rely on GPS/ Global Navigation Satellite System (GNSS), satellite communications, and geospatial data.
- Launch 2-3 low-cost pilots with business units to test the value of commercial satellite data in a specific domain (e.g., supply chain visibility, climate risk monitoring).

DECISION: Who is the executive sponsor for our space strategy?

Near

INTEGRATE & BUILD RESILIENCE

- Integrate proven data sources into core enterprise systems. Move beyond dashboards to embedding insights into automated workflows.
- Design resilience by diversifying providers and architecting hybrid terrestrial/ space networks.

DECISION: Which successful pilots should we scale into operational capabilities?

Next

INNOVATE BUSINESS MODELS

- Explore new business models enabled by next-generation capabilities like orbital compute and ubiquitous connectivity.
- Form strategic partnerships to influence the development of future space-enabled services tailored to your industry.

DECISION: Where can we use space to create a durable, defensible competitive advantage?

ENABLING TECHNOLOGIES

Satellite Connectivity

Mature technology for network resilience

Earth Observation (EO) Analytics

High-ROI data for existing analytics platforms

Persistent Sensing

Evolving capability for continuous monitoring

On-Orbit Edge AI

Emerging area for faster insights

Space-Based Cybersecurity

Capabilities to protect critical space assets

Orbital Compute

Begin small-scale experiments with in-orbit data centers

In-Space Logistics & Manufacturing

Monitor the development of the true in-orbit economy

The Unmanaged Risk: What Happens When Space Fails?

Boards are conditioned to think in failure modes. The current dialogue around space is often focused on opportunity, understating the clear and present operational risk. The resilience of the enterprise is now directly linked to the resilience of space-based systems.

Leaders must ask: What breaks if space systems fail?

- **GNSS/GPS Outage or Degradation:** Precision timing is required for financial transactions, network synchronization, and electrical grids. A loss of signal could disrupt or halt these systems. Logistics and transportation networks that rely on precise navigation would grind to a halt.
- **Satellite Communications (Satcom) Disruption:** Remote industrial operations, maritime fleets, and aircraft rely on satcom for connectivity. An outage would create operational blackouts, posing safety and revenue risks.
- **Earth Observation (EO) Data Interruption:** Companies using EO data for agricultural forecasting, insurance claim verification, and climate risk modeling would lose critical intelligence, leading to poor strategic and financial decisions.

Leadership Implications

Space capabilities are becoming embedded in the systems that support enterprise operations, from connectivity and timing to geospatial intelligence. As reliance grows, space is emerging as a core infrastructure layer shaping performance, resilience, and growth. The task for leadership is to translate this dependence into clear strategic priorities, operating models, and technology decisions.

CEO & Boards - What This Means

Space is no longer a sector decision. It is a core enterprise infrastructure decision.

Enterprise Resilience is Now a Competitive Advantage

- Your operations, supply chains, pricing models, risk analytics, and customer experience increasingly depend on space-enabled systems (connectivity, timing, sensing, geospatial intelligence).
- Space failures are not “technology outages”; they are **enterprise-level disruptions** with revenue, safety, and reputational consequences.
- Leaders who proactively design resilience (multi-orbit, multi-provider, hybrid terrestrial/space architectures) will outperform peers during shocks.

IMPLICATION

Space dependency must be governed like energy, cloud, or cybersecurity, not delegated as an R&D topic.

Space Is Becoming a Growth Platform, Not Just a Cost Center

- Most of the value in the space economy is no longer launch or hardware; it is **data, intelligence, and decision advantage** embedded into operations.
- Companies that integrate space-derived data into products, pricing, underwriting, logistics, and services are creating new revenue streams and margin expansion, not just operational efficiency.

IMPLICATION

Space participation is a growth and differentiation lever, not an exploratory science bet.

Government Decisions Will Shape the Market

- Government agencies are major buyers of space services, so public-sector spending can strongly influence which providers grow and which capabilities scale.
- The same space systems may serve both commercial and national security needs, which means policy, export rules, and procurement decisions can directly affect market access and partnerships.
- For business leaders, that creates opportunity, but it also means greater scrutiny around compliance, reputation, and who you choose to work with.

IMPLICATION

This is not just a technology or vendor decision. Leaders need to understand how regulation, government demand, and geopolitical priorities could affect growth, partnerships, and risk.

THE CEO QUESTION: Where is our business already dependent on space and what happens to revenue, customers, and operations if that dependency fails or becomes constrained?

What This Means for the CTO & CIO

Space is becoming a native extension of the enterprise technology stack.

Space Is Now Part of the Digital Architecture

- Satellite connectivity, Earth observation, and Positioning, Navigation, and Timing (PNT) are no longer “external data sources”—they are becoming **inputs into core systems** (ERP, risk engines, digital twins, AI models).
- The future architecture is **cloud + edge + space**, not cloud alone.

IMPLICATION

Space dependency must be governed like energy, cloud, or cybersecurity, not delegated as an R&D topic.

Data Gravity Is Moving Upward

- As satellite constellations generate continuous, high-volume data streams, pushing all processing to Earth creates latency, bandwidth, and resilience constraints.
- Early orbital compute and edge-AI experiments signal a shift toward **processing data closer to where it is generated**.

IMPLICATION

Space dependency must be governed like energy, cloud, or cybersecurity not delegated as an R&D topic.

Space Expands the Cyber and Risk Surface

- Space systems introduce new attack vectors: RF interference, jamming, spoofing, and orbital congestion.
- Cybersecurity, zero-trust architectures, identity, and monitoring must extend across **ground, link, and orbital segments**.

IMPLICATION

Space assets must be treated as first-class citizens in cyber, continuity, and incident-response planning.

Integration, Not Ownership, Is the Default Model

- Most enterprises will not own satellites—but they will consume space as a service through APIs, platforms, and partnerships.
- The challenge is not access; it is **integration, reliability, and governance**.

IMPLICATION

Success depends on integration capability, vendor orchestration, and architectural discipline—not aerospace expertise.

THE CTO/CIO QUESTION: Which enterprise decisions would materially improve if we had persistent, trusted, space-derived data embedded directly into our systems?

Explore our latest insights


KPMG Cosmos: [Space is Scaling Up](#) | KPMG Cosmos: [Space takes hold](#) | LinkedIn Newsletter: [The Edge of Everywhere](#)

Space
eBooks:



FRONTIER JOURNEY
Building your Blueprint

Download the **Space Economy Blueprint**
to help navigate your journey.



Areas we're watching: Bioeconomy

From refining materials to engineering biology

What if the next industrial revolution is not built on silicon, but biology?

For most of the industrial age, value came from transforming physical materials, refining, forging, and assembling. The digital era moved the frontier to information, and we learned to program it and scale it at near-zero marginal cost. The emerging shift is stranger and more consequential – learning to program biology itself.

What makes this a now story rather than a perennial promise is a collision of cost curves. Reading DNA has fallen faster than Moore's Law for over fifteen years, and AI has solved the design bottleneck — predicting and even inventing functional proteins on demand. Cheap to read, cheap to write, and now designable. Biology has become an engineering substrate.

The implication for production is profound. A living cell, suitably programmed, is a factory: it can manufacture fuels, materials, food, and medicines from renewable inputs. Precision fermentation, cultivated proteins, and bio-based chemicals are already commercial, and much of today's chemical inputs could eventually be made biologically. The unsolved problem isn't the science; it's scaling a process from the lab bench to commodity economics. Whoever cracks that first wins the platform.

This has become a contest between nations, not just companies. Major economies have elevated biomanufacturing to the level of strategic infrastructure, with coordinated industrial policies, multi-billion-dollar investment programs, and explicit positioning around bio-based production as a national priority. The materials economy was won with steel and silicon. The next one may be won with cells, and the positioning is happening now.

Why now?

Three forces are converging to make the bioeconomy a board-level issue:

- 1. Resource security is becoming a strategic priority**

Conflicts in key regions have increased volatility in energy and agricultural markets, exposing how dependent companies are on global supply chains.¹ This is expanding the focus from energy security to broader materials security. Bio-based production offers a way to localize critical inputs and reduce reliance on fragile global systems.^{2,3}

- 2. Global supply chains are being reassessed**

Reliance on overseas manufacturing is now seen as a risk across industries.^{4,5} Disruptions from pandemics, geopolitical tension, and trade constraints are forcing companies to rethink where and how they produce.^{5,6}

- 3. Policy is shaping the competitive landscape**

Governments are using industrial policy to influence where innovation and production happen. Incentives such as the CHIPS and Science Act and the Inflation Reduction Act create opportunity, but also increase complexity and require more active strategic alignment.⁷

Why it matters?

For multinational companies, the operating environment is becoming more complex and less predictable. Bio-based approaches are emerging as a practical way to address core business challenges.

- Supply chain resilience**

Biotechnology enables more localized production of critical inputs. This includes improving processes to recover rare earth elements⁸ and producing biopolymers and other advanced materials domestically⁹, reducing reliance on concentrated global supply sources.⁵

- **Cost stability**

Bio-based inputs can reduce exposure to unpredictable petroleum-based raw material costs and support more stable cost structures. However, broader adoption will depend on improved cost competitiveness, as bio-based equivalents must reach cost neutrality within formulations to achieve meaningful market penetration.⁹

- **Competitive advantage**

Companies can develop proprietary, high-performance materials and products that are difficult for competitors to replicate.¹⁰

- **Alignment with national priorities**

Government incentives, such as the Inflation Reduction Act, can help offset investment costs while aligning corporate strategy with national industrial priorities. Recent legislative changes have modified some tax credits while maintaining support for key clean energy and manufacturing initiatives.

What's next

The science has delivered on its promise. Biology can produce a vast array of useful molecules and materials. The market, however, has not kept pace. This is not a failure of technology but a classic market failure. The landscape is dominated by incumbent industries with decades-old economies of scale and entrenched infrastructure. As a result, despite their technological readiness, novel bio-based products remain a sliver of their potential markets: bioplastics at just 0.5% of global plastics production, biofuels stalled at less than 3% of transport fuel, and bio-based chemicals at an estimated 2%.¹¹ Private sector investment at scale hesitates, unwilling to take on the risk of competing with incumbents without stronger policy support.

For most enterprises, the appropriate posture today is awareness and selective preparation rather than large capital commitment.

In the near term, the practical starting point is understanding where the business is exposed. That means auditing supply chains for critical dependencies on materials or inputs that bio-based alternatives could eventually address, and monitoring which pre-commercial biotechnologies — biomining for critical minerals, biorecovery of precious metals, bio-based polymer substitutes — are advancing toward pilot scale.

As the technology matures, companies with identified vulnerabilities may consider joining industry consortia or forming limited joint ventures with bio-industrial leaders to participate in early pilots without bearing full commercialization risk. Government incentives, where available, can help de-risk these positions.

Longer term, the companies best positioned will be those that have built the relationships, regulatory knowledge, and operational familiarity to move quickly if a bio-based alternative crosses the cost-competitiveness threshold. The valley of death between laboratory readiness and commercial scale is real — but it has been crossed before in materials and energy. The question for leaders is not whether to build a bioeconomy strategy today, but whether they will be ready when the market catches up to the science.

Areas we're watching: The Perception Layer

From blind spots to continuous awareness

Advances in sensing, photonics, and next-generation semiconductors are making the physical world increasingly visible in real time, creating what we call the Perception Layer. It's the capability that turns operations from something periodically measured into something continuously observed.

Why now?

As sensing becomes embedded across operations, it is driving a fundamental upgrade in how systems operate.

- **Reactive » to Real-Time**
- **Estimated » to Observed**
- **Partial Visibility » to Continuous Awareness**

But what does this mean in practice?

- **Logistics:** From reacting to delays to continuously tracking location and condition.¹²
- **Energy and utilities:** From forecasting demand to adjusting consumption in real time.
- **Manufacturing:** From scheduled maintenance to detecting failure signals as they emerge.¹³

This shift enables faster decisions, more reliable automation, and tighter control over complex systems.

The shift

The shift is not simply toward more sensors or faster infrastructure. It is toward systems that can observe, interpret, and respond to the physical world continuously rather than intermittently. That shift is becoming possible because three underlying constraints are moving at once.

- **Real-world data is becoming strategically valuable**
In some applications, AI system performance depends on access to current, real-world data. This is increasing the value of sensing, optical connectivity, and related technologies that can capture and move physical world data at speed and scale.
- **Semiconductor infrastructure is growing to support AI**
AI is fueling major investments in semiconductor systems. High Bandwidth Memory, which AI systems depend on, is growing rapidly.¹⁴ The broader semiconductor materials industry is expanding significantly, projected to grow from \$74.85 billion in 2026 to \$104.22 billion by 2034.¹⁵

At the same time, advanced semiconductors like silicon carbide and gallium nitride are becoming more important for power-efficient systems across industries.¹⁶

- **Energy efficiency is becoming the gating factor for scale**

As AI and digital infrastructure expand, power is becoming a hard constraint. This raises the strategic importance of more efficient semiconductors, architectures, and system design, especially in environments where perception must be continuous and always on.¹⁷

A foundational layer for the frontiers

The Perception Layer is not a standalone trend. It is an enabling layer that supports multiple frontiers and shapes how they scale.

Path to AGI

More advanced AI systems require real-world data to move beyond text-based knowledge and reason about physical environments. The Perception Layer provides that input.

Physical AI

Robots and autonomous systems depend on real-world awareness. The ability to sense and interpret environments becomes a core differentiator.

Energy & Compute

Real-time sensing enables more efficient and resilient infrastructure, supporting dynamic energy optimization, predictive maintenance, and load balancing.

Quantum & NextGen Computing

Advanced systems require precise control and monitoring. Sensing plays a critical role in calibration, stability, and error management.

Future of Trust

Trust depends on a reliable record of what occurred. The Perception Layer can provide a verifiable stream of real-world data to support audit, compliance, and accountability.

Space

Distributed, real-time observation supports autonomous navigation, coordination across assets, and continuous monitoring of space and Earth environments.

What's next?

While the path to AGI remains uncertain, the direction is not. AI will reshape every frontier that matters to business.



As capability advances, diffuses through the economy, and becomes more foundational, leaders will need to rethink strategy, risk, trust, and transformation together.

No area of the business will be untouched. From digital trust and physical AI to quantum and space, the organizations that move early and adapt as conditions change, will be best positioned to shape the next decade rather than just survive it.

The frontiers in this report are not separate stories. They are multiplying, and the leaders who treat them as one strategic landscape rather than six parallel initiatives will set the pace for everyone else.

This report is a starting point, not a conclusion. The signals will keep shifting, and so will our analysis. KPMG is ready to help you turn these shifts into strategy, action, and advantage.

How KPMG can help

KPMG helps organizations turn these frontier shifts into practical advantage, translating emerging signals into strategy, operating choices, and measurable action across trust, transformation, and growth. If you would like to explore what this means for your business, contact one of our KPMG professionals.

Contact a KPMG professional to discuss how we can support your innovation journey.

Appendix

Quantum & NextGen Computing

Definitions and Key Terms

A plain-language reference for the acronyms and technologies used throughout this report

Term	Plain-English Definition
Analog in-memory computing (AIMC)	Performing AI calculations directly inside the memory array in the analog domain, cutting the energy spent moving data to a separate processor.
CHIPS Act	A 2022 U.S. law that funds domestic semiconductor manufacturing and research.
CISA	The Cybersecurity and Infrastructure Security Agency, the U.S. government’s lead body for cybersecurity guidance.
CXL (Compute Express Link)	An open industry standard that lets different types of chips share memory and work together in one system.
DARPA	The Defense Advanced Research Projects Agency, the U.S. Defense Department’s advanced-research arm.
DNA data storage	Storing digital information in synthetic DNA molecules for extremely dense, long-lasting archival storage.
DOE	The U.S. Department of Energy, a major funder of national computing and quantum research.
Energy-based model (EBM)	A type of AI model that scores each possible outcome by an “energy” value, where lower energy means more likely.
Error correction / fault tolerance	Techniques that detect and fix the errors quantum hardware makes; required for large, reliable quantum computers.
Exascale computing	Supercomputers able to perform at least a billion-billion calculations per second.
Harvest now, decrypt later	The risk that attackers collect encrypted data today in order to unlock it once quantum computers can break the encryption.
High-bandwidth memory (HBM)	Fast memory stacked close to a processor, widely used in AI hardware.
High-performance computing (HPC)	Large-scale computing using supercomputers or clusters for the most demanding workloads.
Hybrid architecture	A design in which quantum, classical, and AI systems work together, each handling the part of a problem it does best.
Moore’s Law	The long-running trend of chip performance roughly doubling every couple of years, now delivering diminishing economic returns.

Neuromorphic computing	Chips modeled on the brain that process information in an energy-efficient, event-driven way.
NIST	The National Institute of Standards and Technology, the U.S. agency that sets technical standards, including the new post-quantum encryption standards.
NQI (National Quantum Initiative)	The U.S. national program coordinating quantum research, funding, and workforce development.
OMB	The Office of Management and Budget, the White House office that issues federal management and budget directives.
p-bit (probabilistic bit)	A hardware element that produces controlled randomness, used in thermodynamic and probabilistic chips.
Phase-change memory (PCM)	A memory material that stores data by switching physical states; used in analog in-memory chips.
Photonic computing	Using light instead of electricity to move and process data, improving speed and energy efficiency.
Post-quantum cryptography (PQC)	New encryption methods designed to stay secure against future quantum computers.
Probabilistic and approximate computing	A family of approaches that trade exact precision for speed and energy savings where perfect accuracy is not needed.
Processing-in-memory (PIM)	Performing computation inside or beside memory to reduce costly data movement.
Quantum advantage	The point at which a quantum computer solves a useful problem better or faster than the best classical computer.
Quantum-centric supercomputing	Combining quantum processors with classical supercomputers in one integrated system.
Quantum sensing	Using quantum effects to measure motion, time, gravity, or magnetic fields with extreme precision.
Qubit	The basic unit of quantum information, the quantum counterpart to a classical 0/1 bit.
Reversible computing (Near-Zero Energy Chips)	Chips that recycle the energy of each computation instead of releasing it as heat, dramatically reducing power use.
RIKEN	Japan's national research institute, a partner in recent quantum-supercomputing work.
RRAM (Resistive RAM)	A memory technology used in analog in-memory computing.
Superconducting / cryogenic computing	Computing with materials cooled to near absolute zero, where electrical resistance vanishes, for very high energy efficiency.
Thermodynamic computing	Chips that harness natural electrical noise to perform AI sampling tasks at far lower energy than conventional processors.
TSU (Thermodynamic Sampling Unit)	The specialized chip the startup Extropic builds for thermodynamic computing.

Our Methodology

In examining the forces that will shape our future, we begin not with assumptions but with questions. Where will the initial signs of change appear? Which patterns are emerging, speeding up, or collapsing? What fundamental dynamics are altering the competitive landscape in ways that are most significant for decision-makers now and in the future?

We use a mixed-method approach that integrates human intuition with data analysis, monitoring for initial signals of disruption and organizing them into patterns that uncover deeper insights into market trajectories. These signals could include headlines, product announcements, expert opinion, client sentiments, capital flows, and other quantitative and qualitative data points that ultimately lay the groundwork for transformative trends.

Each signal we identify is carefully evaluated across four essential dimensions:

- Strategic relevance to long-term business goals
- Ability to disrupt or redefine industries
- Readiness for adoption considering existing barriers and enablers
- Speed at which the signal is transitioning from niche to mainstream

From these signals, we then extract macro trends—the significant movements gaining momentum across various sectors. We evaluate trends by applying a set of methodologies to highlight those that promise to be the most impactful.

Foresight gains its value when it encourages action, which is why we include grounded recommendations for leaders across both short and long-term time horizons. Our signals-based strategy model can help organizations prioritize innovation, mitigate risks, and uncover new opportunities. Designed to be applicable across various sectors—from energy and finance to healthcare and technology—these tools reflect the fact that the dynamics of the future will transcend industry boundaries.

Sources

Path to AGI

1. METR, "Measuring AI Ability to Complete Long Tasks," March 2025.
2. DeepMind, "AlphaEvolve: A coding agent for scientific and algorithmic discovery," White Paper, May 2025.
3. International Organization for Standardization, ISO/IEC 42001:2023 — Artificial Intelligence Management System, 2023.
4. National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0), 2023.
5. European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act).
6. Okta, "AI at Work 2025: Securing the AI-Powered Workforce," 2025.
7. Microsoft, "Cyber Pulse: AI Security Report," February 2026.
8. Cloud Security Alliance and Strata Identity, "Securing Autonomous AI Agents," February 2026.
9. KPMG, "2025 CEO Outlook," 2025.
10. RAND Corporation, "Liability for Harms from AI Systems: The Application of U.S. Tort Law," 2024.
11. KPMG, "AI Quarterly Pulse Survey Q2 2026," June 2026.
12. Stanford Digital Economy Lab, "The Effects of AI on Productivity and Work Practices," (n.d.).
13. Financial Stability Board, "The Financial Stability Implications of Artificial Intelligence," November 2024.
14. CNBC, "JPMorgan Chase is giving its employees an AI assistant powered by ChatGPT maker OpenAI," August 2024; PYMNTS, "Goldman Sachs Makes AI the Centerpiece of Q3 Earnings," October 2025.
15. TechCrunch, "Shopify CEO tells teams to consider using AI before growing headcount," April 2025.

Physical AI

1. Mind the Bridge, "Silicon Valley bets big on physical AI: 93% of VC flows now go into artificial intelligence," September 2025.
2. Reuters, "ABB to sell robots business in \$5.38 billion deal involving SoftBank," October 2025.
3. DeepMind, "RT-2: New model translates vision and language into action," Google DeepMind, 2023.
4. Fortune Business Insights, "Edge AI market size, share & industry analysis," 2026.
5. The Insight Partners, "Digital Twin Market Forecast to 2031," 2025.
6. European Commission, Draft Guidelines on the Classification of High-Risk AI Systems under the AI Act, Article 6, May 2026.
7. Mobility Global, "How autonomous vehicles will change the future of car insurance," 2025.
8. National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0), 2023.
9. International Federation of Robotics, "World Robotics 2025," 2025.
10. The Robot Report, "From EVs to robotics: Tesla targets 10M Optimus units with new Texas plant," April 2026.
11. Amazon News, "Amazon has more than 1 million robots that sort, lift, and carry packages—see them in action," 2025.
12. Meta AI, "Advancing embodied AI through progress in touch perception, dexterity, and human-robot interaction," Meta FAIR, October 2024.
13. Hyundai Motor Group, "Hyundai Motor Group Launches Boston Dynamics AI Institute," August 2022.

Energy & Compute

1. International Energy Agency, "Energy and AI," 2025.
2. Brookings, "Eight Facts About Permitting and the Clean Energy Transition," May 2024.
3. U.S. Department of Energy, "Coordinated Interagency Transmission Authorizations and Permits Program."
4. Micron Technology Inc., "Fiscal Q1 2026 Earnings Call Prepared Remarks," December 2025.
5. PJM Interconnection, "2027/2028 Base Residual Auction Report," December 2025.
6. CNBC, "Top Hyperscalers to Boost AI Capex to \$600 Billion: Stocks That Benefit," February 2026.
7. M&G Investments, "Tech Issues: The AI Debt Deluge Hitting Bond Markets," March 2026.
8. ConstructConnect, "October Data Center Report: Spending Reaches \$13B as Costs Rise," October 2025.

9. Edison Electric Institute, "2024 Financial Review," July 2025.
10. U.S. Energy Information Administration, "U.S. Electricity Prices Continue Steady Increase," May 2025.
11. The Wall Street Journal, "SpaceX and Google Are in Talks to Launch Data Centers in Orbit," May 2026.
12. PNAS, "Can Neuromorphic Computing Help Reduce AI's High Energy Cost?," October 2025.
13. Microsoft, "AI Chips Are Getting Hotter. A Microfluidics Breakthrough Goes Straight to the Silicon to Cool Up to Three Times Better," January 2026.
14. Fusion Industry Association, "The Global Fusion Industry in 2025," July 2025.
15. U.S. Department of Energy, "Energy Department Selects TVA and Holtec to Advance Deployment of U.S. Small Modular Reactors," December 2025.
16. Reuters, "Cyberattacks on US Utilities Surged 70% This Year, Says Check Point," September 2024.
17. CISA, "Iranian-Affiliated Cyber Actors Exploit Programmable Logic Controllers Across US Critical Infrastructure," April 2026.
18. KJZZ, "Developers Push to Close Proposed Project Blue Data Center Land Deal by Christmas," November 2025.
19. Data Center Dynamics, "Judge Rules Zoning for PW Digital Gateway Data Center Campus in Virginia Voided," August 2025.
20. The New York Times, "A Mysterious Company Came to Town With a \$165 Billion Idea," December 2025.
21. EPRI, "DCFlex Initiative," 2024–present.
22. Duke Energy, "Responding to Growing Demand, Duke Energy, Amazon, Google, Microsoft and Nucor Execute Agreements to Accelerate Clean Energy Options," May 2024.
23. CNN, "Elon Musk Is Building 'the World's Biggest Supercomputer.' It's Powered with Dozens of Gas-Powered Turbines," May 2025.
24. PV Tech, "Amazon Buys 1.2GW Sunstone Solar-Plus-Storage Site," January 2026.
25. Federal Energy Regulatory Commission, "Ensuring the Timely and Orderly Interconnection of Large Loads," Advance Notice of Proposed Rulemaking, October 2025.
26. Federal Energy Regulatory Commission, "FERC Directs Nation's Largest Grid Operator to Create New Rules to Embrace Innovation and Protect Consumers," December 2025.
27. Google Blog, "Google Signs Advanced Nuclear Clean Energy Agreement with Kairos Power," October 2024.
28. Helion Energy, "Helion Announces World's First Fusion Energy Purchase Agreement with Microsoft," May 2023.
29. SAFELOG GmbH, "SAFELOG Relies on Innovative Robot Localisation from Opteran," July 2024.
30. KHOU, "Microgrid Helped H-E-B Stay Open While Most of Houston Was Without Power," August 2024.
31. Alphabet Inc., "Alphabet Announces Agreement to Acquire Intersect to Advance U.S. Energy Innovation," December 2025.
32. Voltus, "Voltus to Deliver \$240 Million to Customers in Record 2025," December 2025.
33. TSMC, "TSMC's EUV Dynamic Power Saving: A Win-Win for Energy Efficiency and Production Capacity," September 2025.

Quantum & NextGen Computing

QUANTUM

1. National Institute of Standards and Technology, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards," August 2024.
2. Wall Street Journal, "U.S. to Award Quantum-Computing Firms \$2 Billion and Take Equity Stakes," May 2026.
3. PsiQuantum, "PsiQuantum Signs \$100M Letter of Intent with U.S. Department of Commerce," May 2026.
4. Federal Ministry of Education and Research, "Quantum Technologies Conceptual Framework Programme."
5. European Commission, "European Quantum Communication Infrastructure – EuroQCI | Shaping Europe's Digital Future."
6. UK Research and Innovation, "NQTP | UK National Quantum Technologies Programme."
7. Qureca, "Quantum Initiatives Worldwide 2025," 2025.
8. Office of Management and Budget, "Memorandum M-23-02: Migrating to Post-Quantum Cryptography," November 2022.
9. Cybersecurity and Infrastructure Security Agency, "Quantum-Readiness: Migration to Post-Quantum Cryptography," August 2023.
10. Nextgov, "Draft executive order would set deadlines for digital signature and key quantum encryption," May 2026.
11. The Quantum Insider, "Why Timing Affects the Cost of Post-Quantum Migration," May 2026.
12. IBM, "IBM Releases a New Blueprint for Quantum-Centric Supercomputing," March 2026.

13. Q-CTRL, "Q-CTRL Delivers 3,000x Speedup in Materials Discovery for the Energy Sector with Quantum Computing, Demonstrates Evidence of Practical Quantum Advantage," May 2026.
14. IBM, "Cleveland Clinic, RIKEN, and IBM Model a 12,635-Atom Protein – the Largest Known to be Simulated with Quantum Computers," May 2026.
15. HSBC News, "HSBC Demonstrates World's First-Known Quantum-Enabled Algorithmic Trading with IBM," September 2025.
16. U.S. Department of Energy, Office of Science, "Quantum Testbed Pathfinder," 2023.
17. Defense Innovation Unit, "DIU to Kick Off Quantum Sensor Field Testing Under TQS."
18. U.S. Department of Energy, "Quantum Information Science Applications Roadmap," 2024.
19. DARPA, "DARPA Advances Quantum Benchmarking Initiative to Stage B with 11 Companies for Utility-Scale Quantum Computing by 2033," November 2025.
20. U.S. Department of Energy, "Energy Department Announces \$625 Million to Advance the Next Phase of National Quantum Information Science Research Centers," November 2025.
21. National Science and Technology Council, "National Quantum Initiative Supplement to the President's FY 2025 Budget," 2024.
22. IBM, "IBM and U.S. Department of Commerce Announce America's First Purpose-Built Quantum Foundry, Supported by Proposed \$1 Billion CHIPS Award," May 2026.
23. Quantum Computing Report, "QuEra, Harvard, and MIT Demonstrate 2:1 Physical-to-Logical Qubit Ratio," April 2026.
24. Nextgov, "OMB Draft Memo Sets Agency and Vendor Quantum Security Standards," July 2025.
25. IBM, "Quantum + AI + HPC: IBM and AMD's Game-Changing Alliance," August 2025.
26. Defense Innovation Unit, "DIU to Kick Off Quantum Sensor Field Testing Under TQS."
27. The Banker, "How ready are banks for quantum's Q-day?," April 2025.
28. Boeing, "Beyond GPS: Team completes 1st quantum navigation flight test," March 2025.

NEXTGEN COMPUTING

1. Cerebras Systems, "Cerebras Systems Announces Pricing of Initial Public Offering," May 2026.
2. Intel Corporation, "Taking Neuromorphic Computing to the Next Level with Loihi 2," Technology Brief, September 2021.
3. Lightmatter, "Lightmatter Achieves Record 1.6 Tbps Per Fiber to Accelerate AI Optical Interconnect," March 2026.
4. IBM Research, "An Energy-Efficient Analog Chip for AI Inference," 2025.
5. Extropic, "An Efficient Probabilistic Hardware Architecture for Diffusion-Like Models," October 2025.
6. Extropic, "Inside X0 and XTR-0," October 2025.
7. EE Times, "Vaire Demos Energy Recovery With Reversible Computing Test Chip," September 2025.
8. imec, "Superconducting Computing Research Project," June 2025.
9. Microsoft Research and University of Washington, "End-to-End Automation of DNA Data Storage," Nature Scientific Reports, March 2019.
10. DNA Data Storage Alliance, "Why DNA Data Storage," (n.d.).
11. Los Alamos National Laboratory, "Neuromorphic computing: the future of AI," March 2025.
12. Oak Ridge Leadership Computing Facility, "Quantum Computing User Program."
13. NeuromorphicCore, "Neuromorphic Computing Investment Trends," June 2025.
14. CNBC, "Nvidia buying AI chip startup Groq's assets for about \$20 billion in its largest deal on record," December 2025.
15. Compute Express Link Consortium, "CXL 2.0 Specification," February 2024.
16. Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics, 2021.
17. NVIDIA, "Siemens and NVIDIA Expand Partnership Industrial AI Operating System," January 2026.
18. Mayo Clinic, "Mayo Clinic Accelerates Personalized Medicine Through Foundation Models with Microsoft Research and Cerebras Systems," January 2025.
19. Advanced Micro Devices (AMD). "5th Gen AMD EPYC™ Processor Architecture." White Paper, October 2024.
20. NVIDIA, "NVIDIA NVLink Fusion Semi-Custom AI Infrastructure Partner Ecosystem," May 2025.
21. 8ra Consortium, "IPCEI-CIS," 2026.
22. Bosch, "Bosch Center for Artificial Intelligence," 2026.

Future of Trust

1. Coalition for Content Provenance and Authenticity, "C2PA Specification," 2024.
2. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework," January 2023.
3. European Union, Regulation (EU) 2024/1689, Artificial Intelligence Act, 2024.
4. International Organization for Standardization, "ISO/IEC 42001: Artificial Intelligence Management Systems," 2023.
5. Pindrop, "2025 Voice Intelligence & Security Report," June 2025.
6. Stanford Institute for Human-Centered AI, "AI Index Report 2025," April 2025.
7. Securitize and BlackRock, "BlackRock's BUIDL, Tokenized by Securitize, Now Accepted as Collateral for Trading on Binance and Launches on BNB Chain," November 2025.
8. DataIntel, "Verifiable Credentials Market Research Report 2023," September 2025.
9. Securitize, "BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network," March 2024.
10. National Institute of Standards and Technology, "Post-Quantum Cryptography Standards: FIPS 203, 204, 205," August 2024.

Space

1. Brian Miske, "The \$1.8 Trillion Question: Making Bioscience in Space Investable," December 2025.
2. Brian Miske, "How AI Turns the Edge of Everywhere Into Value," November 2025.
3. KPMG, "Cosmos Quarterly: Space Economy Q4 2025," February 2026.
4. Brian Miske, "The Next Frontier of Capital: Investing in the Sovereign-Backed Space Economy," October 2025.

Areas we're watching

1. World Economic Forum, "Strategic Intelligence," [ongoing platform].
2. Science Direct, "Biomass to Biofuel: Impacts and Mitigation of Environmental, Health, and Socioeconomic Challenges," January 2025.
3. Lawrence Berkeley National Laboratory, "Bioenergy and Biofuels."
4. Manufacturing Curated, "Can the U.S. Reduce Dependence on Overseas Pharma Manufacturing?" June 2025.
5. The Aspen Institute Economic Strategy Group, "Manufacturing Resilience: The US Drive to Reorder Global Supply Chains," November 2023.
6. Wiley, "Encounter the Unforeseen: Resilient Supply Chain Modeling for a Sustainable Bioeconomy," Journal of Industrial Ecology, 2025.
7. The Harvard Business School - Institute for Business in Global Society, "The Rebirth of U.S. Industrial Policy?" June 2024.
8. UC Berkeley Engineering, "Researchers Pioneer Greener Way to Extract Rare Earth Elements," November 2025.
9. Science Direct, "Ranking Economic and Environmental Performance of Feedstocks Used in Bio-based Production Systems," 2025.
10. Chemical Research Insight, "Top 10 Companies in the Global Sustainable Biopolymers Industry (2025): Market Leaders Driving the Green Materials Revolution," June 2025.
11. European Commission, "Bio-based Plastics: Technology Readiness and High Cost Key Barriers to Wide Use," March 2026.
12. Supply Chain Dive, "Real-time Location and Condition Tracking Drives Resilient Supply Chains," January 2025.
13. Sigma Technology, "Predictive Maintenance Examples from 10 Leading Industries," August 2025.
14. TrendForce, "Memory Wall Bottleneck: AI Compute Sparks Memory Supercycle," January 2026.
15. Fortune Business Insights, "Semiconductor Materials Market Size, Share, and Trends (2026–2034)," 2026.
16. TimesTech, "Wide-bandgap Semiconductor Market Outlook (SiC, GaN)," 2024–2025.
17. RAND Corporation, "AI's Power Requirements Under Exponential Growth," January 2025.