Capgemini

Top Tech Trends of 2026

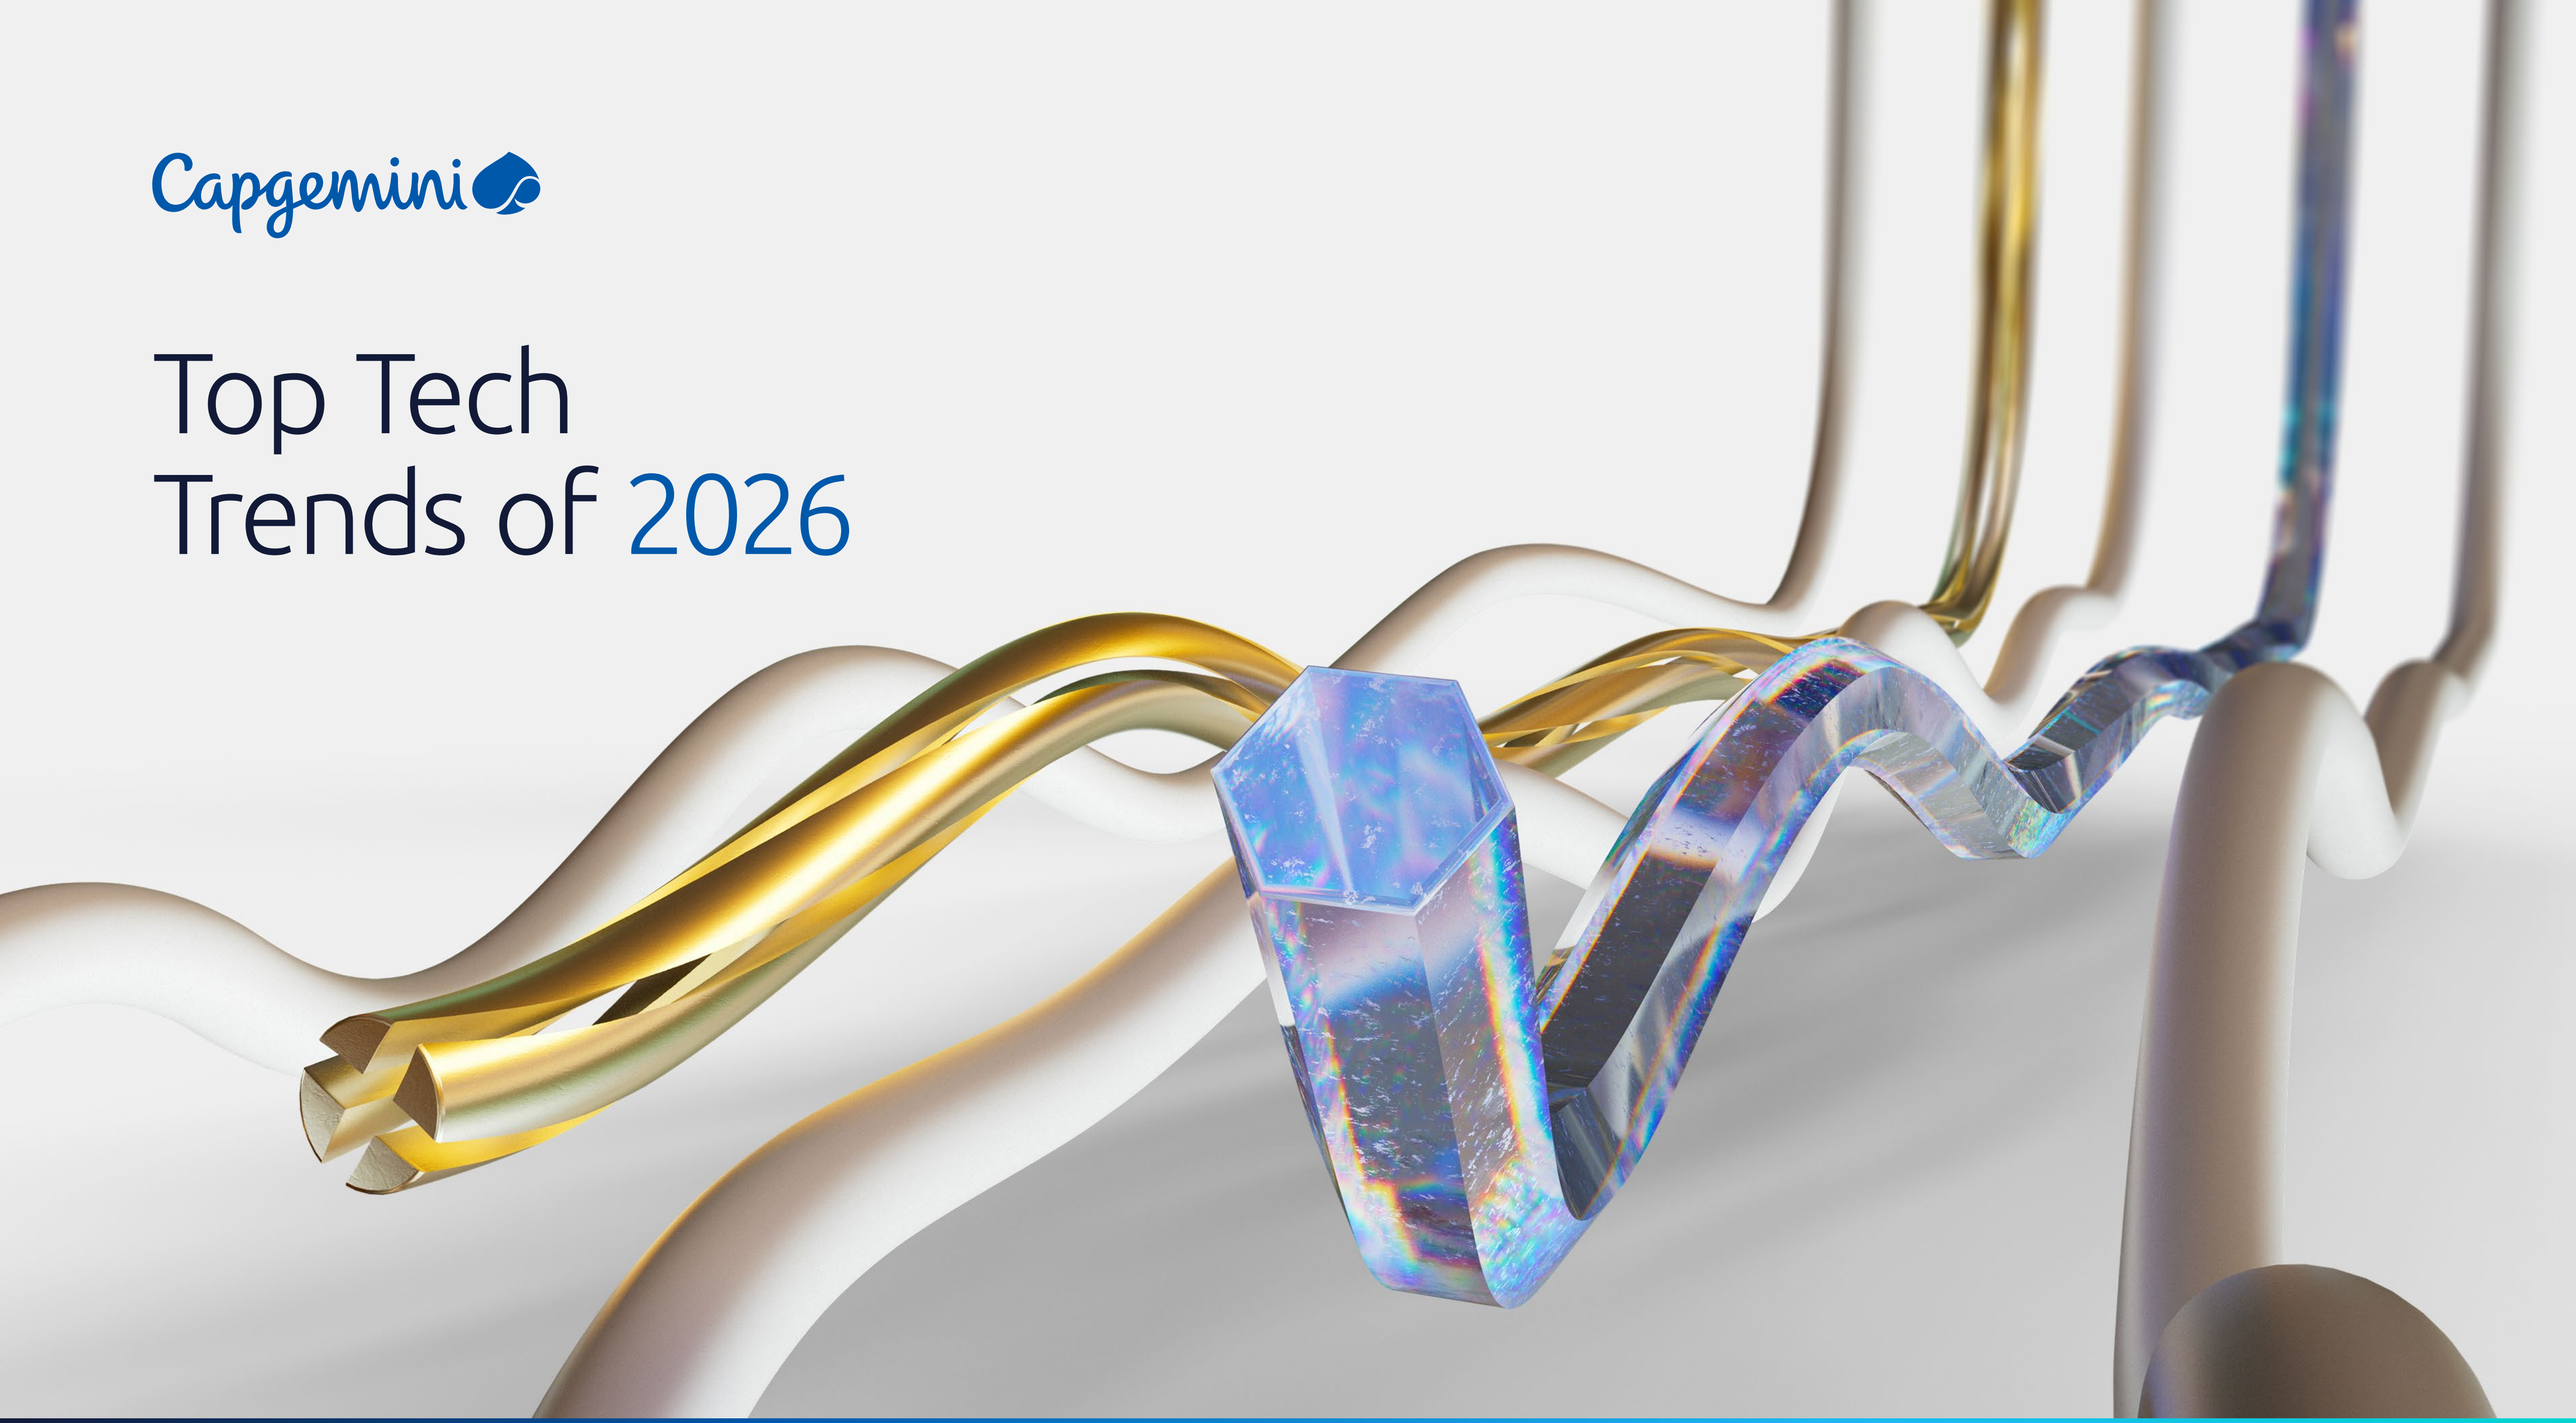

Who should read this report and why?

This report is designed for C-suite executives and business and innovation leaders. The report presents our convictions regarding what will be the most impactful technological trends of 2026. It offers valuable insights into the trends we see dominating the tech landscape, while looking back on the accuracy of our predictions for 2025.

Data from comprehensive surveys of industry executives, the investor community, and in-depth discussions with experts support our predictions. The insights we derive from this analysis will help technology and business leaders in establishing sound strategies and impactful investments.

Table of contents

Pascal Brier

Group Chief Innovation Officer, Member of the Group Executive Committee, Capgemini

Introduction

After several years of extraordinary acceleration across AI, cloud, data, and automation, 2026 marks a shift toward strengthening, upgrading or rebuilding the foundations that will support the next decade.

Across industries, leaders recognize that progress cannot rest on fragmented pilots or loosely connected digital initiatives. The era of experimental AI is giving way to the need for solid AI foundations: reliable data, clear governance, scalable architectures, and systems designed for safety, trust, and measurable outcomes. The organizations able to move from isolated models to integrated, enterprise-wide intelligence will be those that generate lasting value.

At the same time, the global environment is forcing companies to rethink resilience and business continuity at a much deeper level. Rising dependencies on critical technologies (from semiconductors and cloud services to AI models and compute infrastructure) have become strategic risk factors rather than purely technical choices. This is driving a dual movement: a renewed push for architectures that can withstand disruption, and a search for greater control over the layers of technology that matter most. Cloud strategies are evolving accordingly, with hybrid, multi-cloud, and sovereign options emerging not as exceptions but as mechanisms to secure continuity, reduce concentration risk, and safeguard data and operations. Sovereignty is part of this shift, but the underlying theme is broader: organizations are redesigning their foundations to remain open, scalable, and globally connected, while ensuring that no single dependency can compromise their ability to operate.

Our Top Tech Trends for 2026 reflect this shift toward structural rebuilding, pointing to a single message: technology leadership in 2026 is no longer about experimentation, but about constructing the durable foundations that will enable true value to be extracted from innovation.

As every major technological shift has shown, it is the strength of these foundations, not the novelty of individual tools, that determines who captures long-term advantage. This report aims to help business and technology leaders make the right strategic choices at a moment when those foundations are being rebuilt.

Looking back to the top tech trends of the last two years

2025	2024
Generative AI: From copilots to reasoning AI agents	**Generative AI:** Small will be the new big
Cybersecurity: New defenses, new threats	**Quantum technologies:** When cyber meets quantum
AI-driven robotics: Blurring the lines between humans and machines	**Semiconductors:** Moore’s Law isn’t dead, but it is changing
Nuclear: The surge of AI driving the clean tech agenda	**Batteries:** The power of new chemistry
New-generation supply chains: Agile, greener and AI-assisted	**Space tech:** Addressing the Earth’s challenges from outer space

The top tech trends of 2026

The year of truth for AI

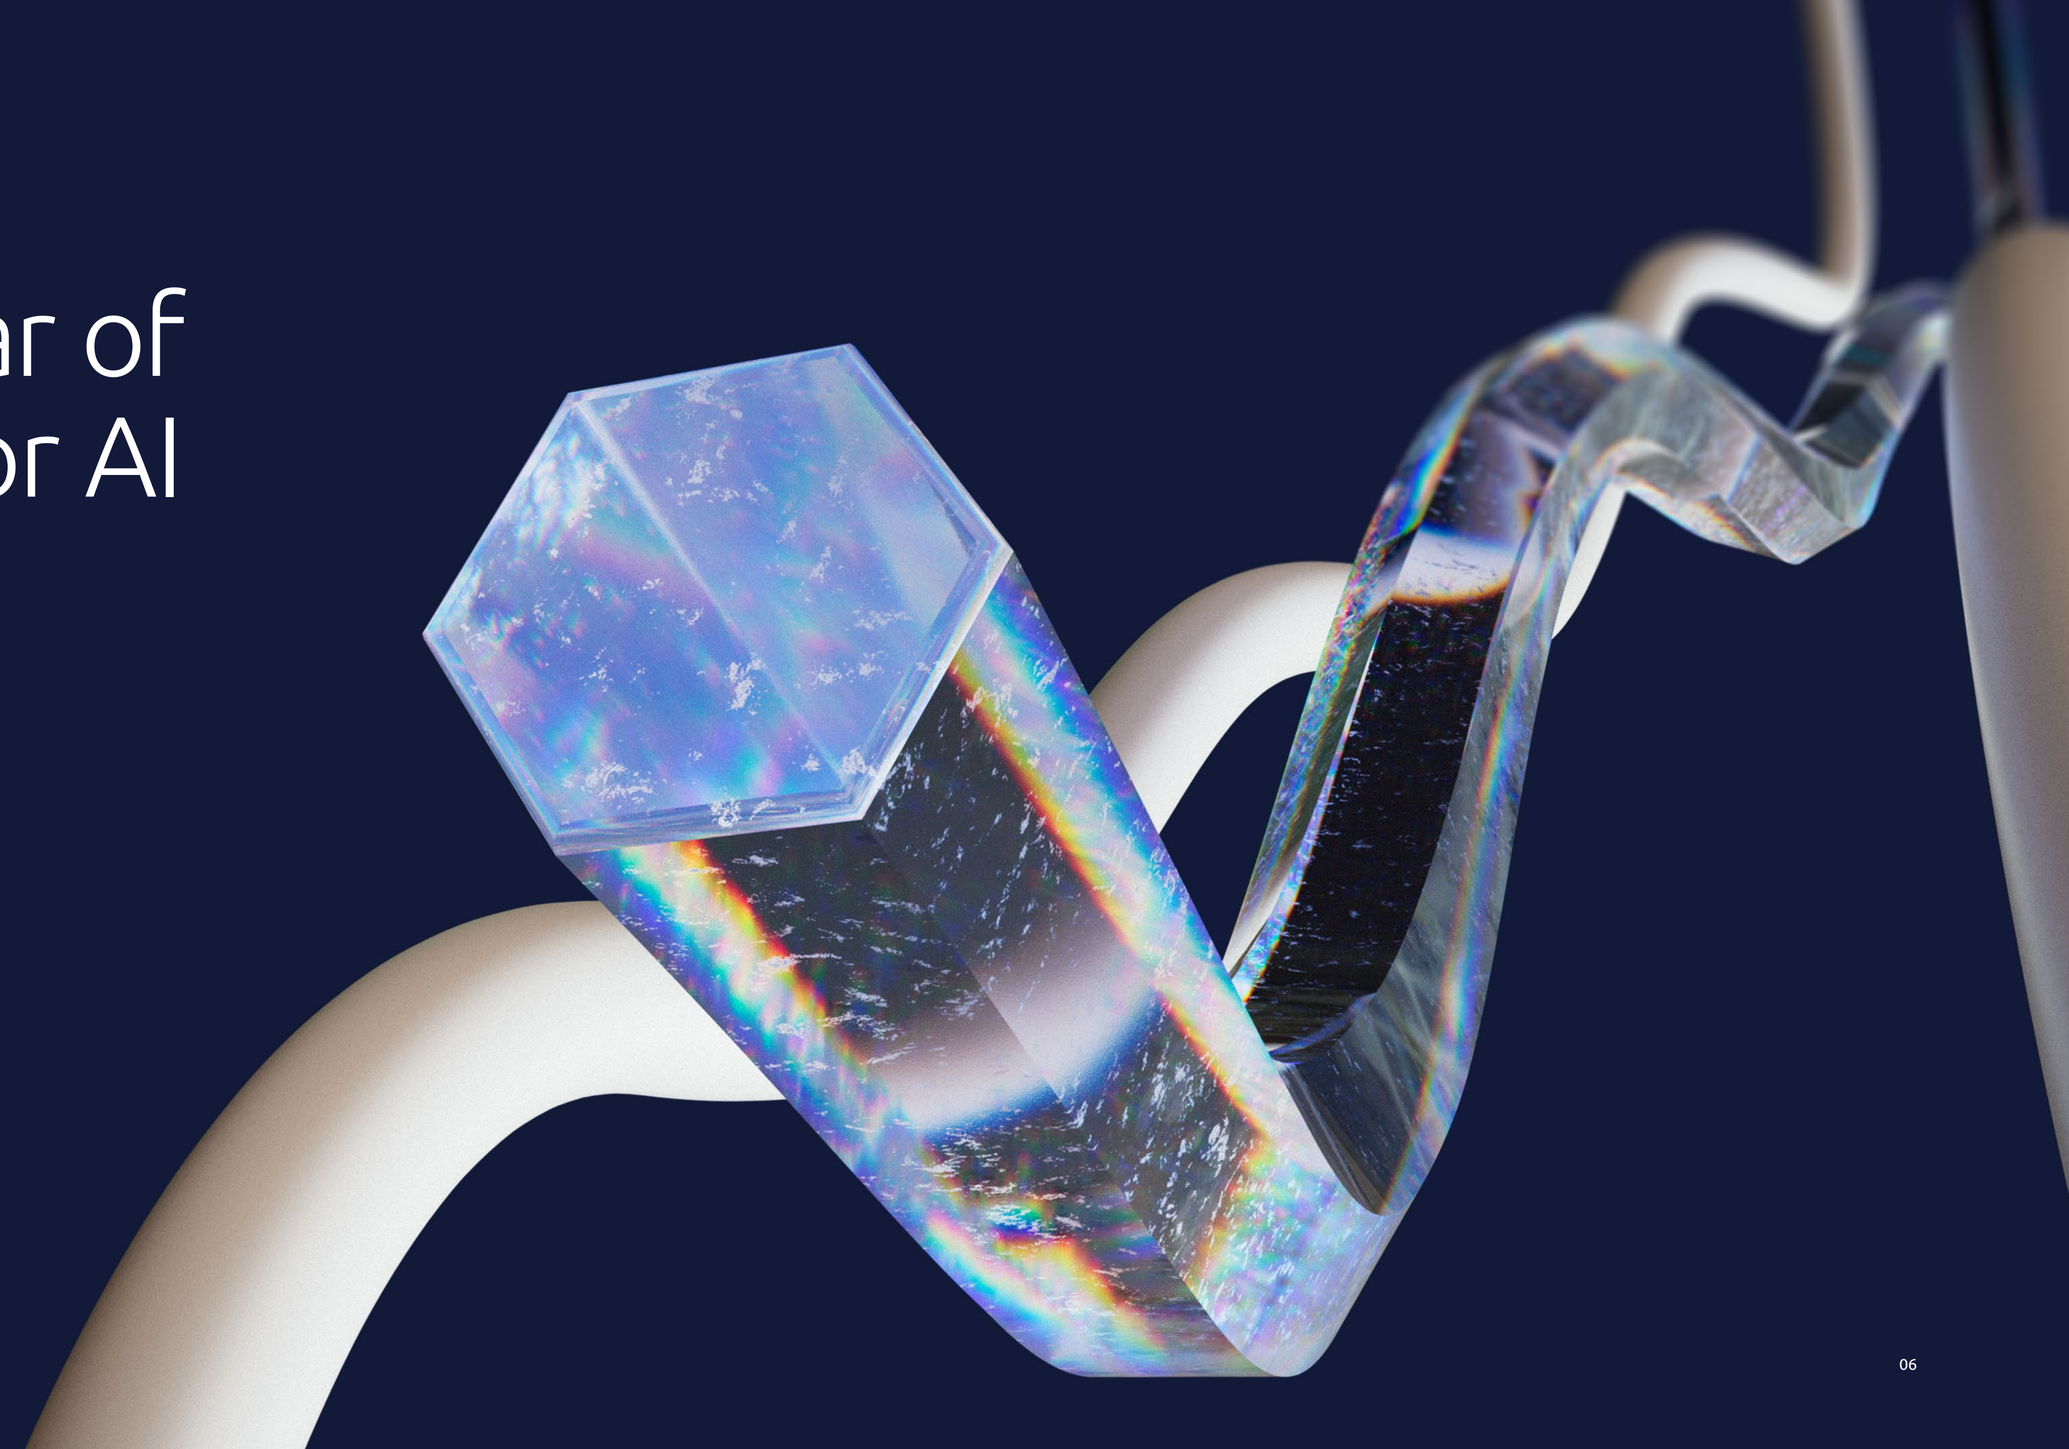

The year of truth for AI

After a period of unprecedented investment and experimentation, AI has become the defining technology of the decade. Yet the pace of investment has outstripped the speed at which organizations have been able to deploy it at scale and extract measurable value. Many enterprises now find themselves with sophisticated models, agents, and prototypes that remain unintegrated, under-utilized, or disconnected from real business outcomes. This gap has generated some skepticism and a sense of some form of AI hype.

Beneath the noise, however, something more consequential is taking shape. The structural foundations of AI are maturing. Organizations that have moved beyond pilots are already seeing tangible results, with early adopters reporting productivity gains of 7–18% across core digital and software operations[1]. Crucially, these gains are not merely absorbed as efficiency: half of organizations reinvest the time saved into developing new features, while nearly as many channel it into workforce upskilling. This marks a shift from experimentation to value compounding.

At the same time, AI itself is evolving in form and function. Large models are becoming more modular, agents are moving from novelty tools to workflow orchestrators, and AI is shifting from peripheral experimentation to deeper integration within enterprise cores. Adoption reflects this transition. Today, roughly 46% of the software workforce uses generative AI tools; by 2026, that figure is expected to reach 85%[2], signaling a move from early adoption to default capability.

This is why 2026 emerges as the year of truth for AI. Short-term hype fades, but what remains is an ecosystem increasingly grounded in operational value, enterprise architecture, and sustained productivity. As with past technology waves, real growth begins once organizations recognize that value does not lie in isolated use cases but in enterprise-wide systems that evolve and scale over time.

Reaching that future requires discipline. Organizations must confront their true AI readiness, starting with data foundations and infrastructure. The agentic wave is accelerating, but not all agents are built to scale; hastily assembled “toy agents” risk renewing disappointment. Differentiation no longer comes from the models themselves, which are rapidly commoditizing, but from architecture, integration, orchestration, and the ability to turn AI into durable, compounding business value.

Why it matters

After years of hype and fragmented pilots, AI can no longer be innovation theater. Investment has outpaced value delivery, and 2026 is the moment when organizations must move from proof-of-concept to proof-of-impact. The next wave of AI is not about specific tools or model releases; it is about embedding intelligence into the fabric of operations, processes, and society (and making it work for everyone). Leaders who succeed will build the capabilities, governance, and human-AI chemistry required to deliver measurable outcomes at scale, while laying the foundations for the larger-scale transformation that will follow.

What to look out for

Organizations increasingly turn from experimental AI agents to **production-grade agentic systems** built to operate within real enterprise architectures. This shift favors platforms that integrate directly with existing data pipelines, identity layers, workflow engines, and business applications. Early proofs of concept give way to robust orchestration frameworks that coordinate multiple specialized agents, each with defined roles, evaluation loops, and governance controls. Momentum also builds around **modular and domain-specific models.**

As general-purpose LLMs commoditize, enterprises prioritize smaller, fine-tuned models tailored to finance, healthcare, retail, or industrial operations. These models rely on improved retrieval, vector databases, and continuous fine-tuning pipelines, giving organizations tighter control over accuracy, provenance, and performance—especially in regulated environments.

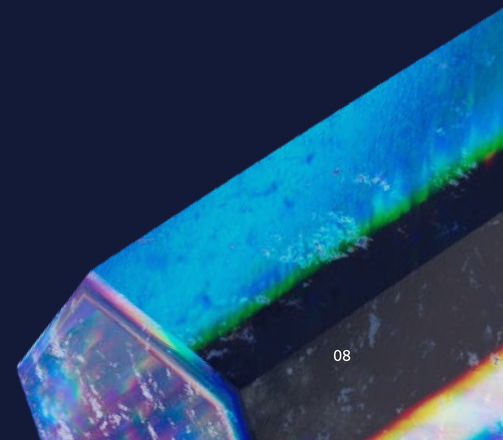

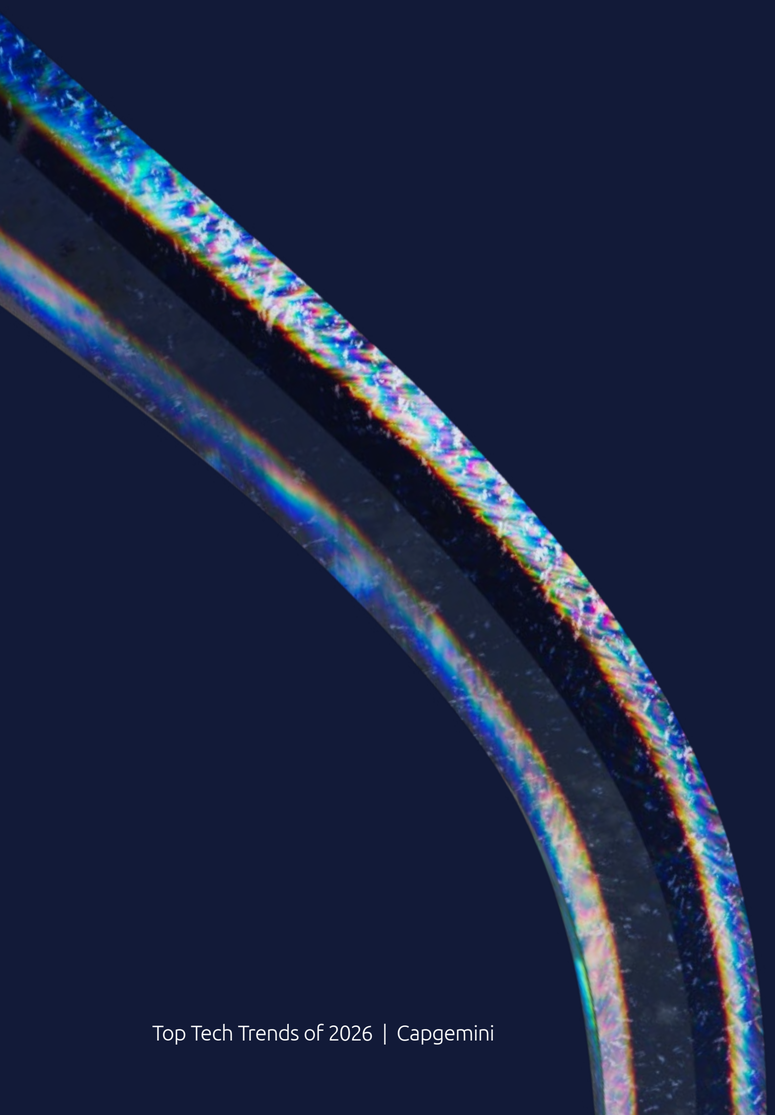

At the same time, the pressure to demonstrate concrete ROI accelerates investment in **AI observability, evaluation, and value measurement.** Companies establish internal evaluation suites to test model behavior, monitor agent decisions, and assess reliability against business outcomes. Dedicated "AI value offices" or governance teams emerge to oversee performance at scale, drawing on telemetry, productivity insights, and financial impact.

Finally, the rapid adoption of **AI-augmented engineering and operations** signals that intelligence is becoming embedded across the software development lifecycle and core business workflows. Code generation pipelines, automated testing agents, self-optimizing data workflows, and AI copilots for operations move from experimentation to standard practice. These shifts reinforce the broader transformation described across this report: AI is no longer an add-on—it is becoming a structural capability of the modern technology stack.

"Generative AI has attracted unprecedented levels of investment and attention. As we move into 2026, the conversation is shifting decisively toward value creation—moving beyond experimentation to measurable business impact."

Mark Roberts

Head of AI Futures Lab, Capgemini

AI is eating software

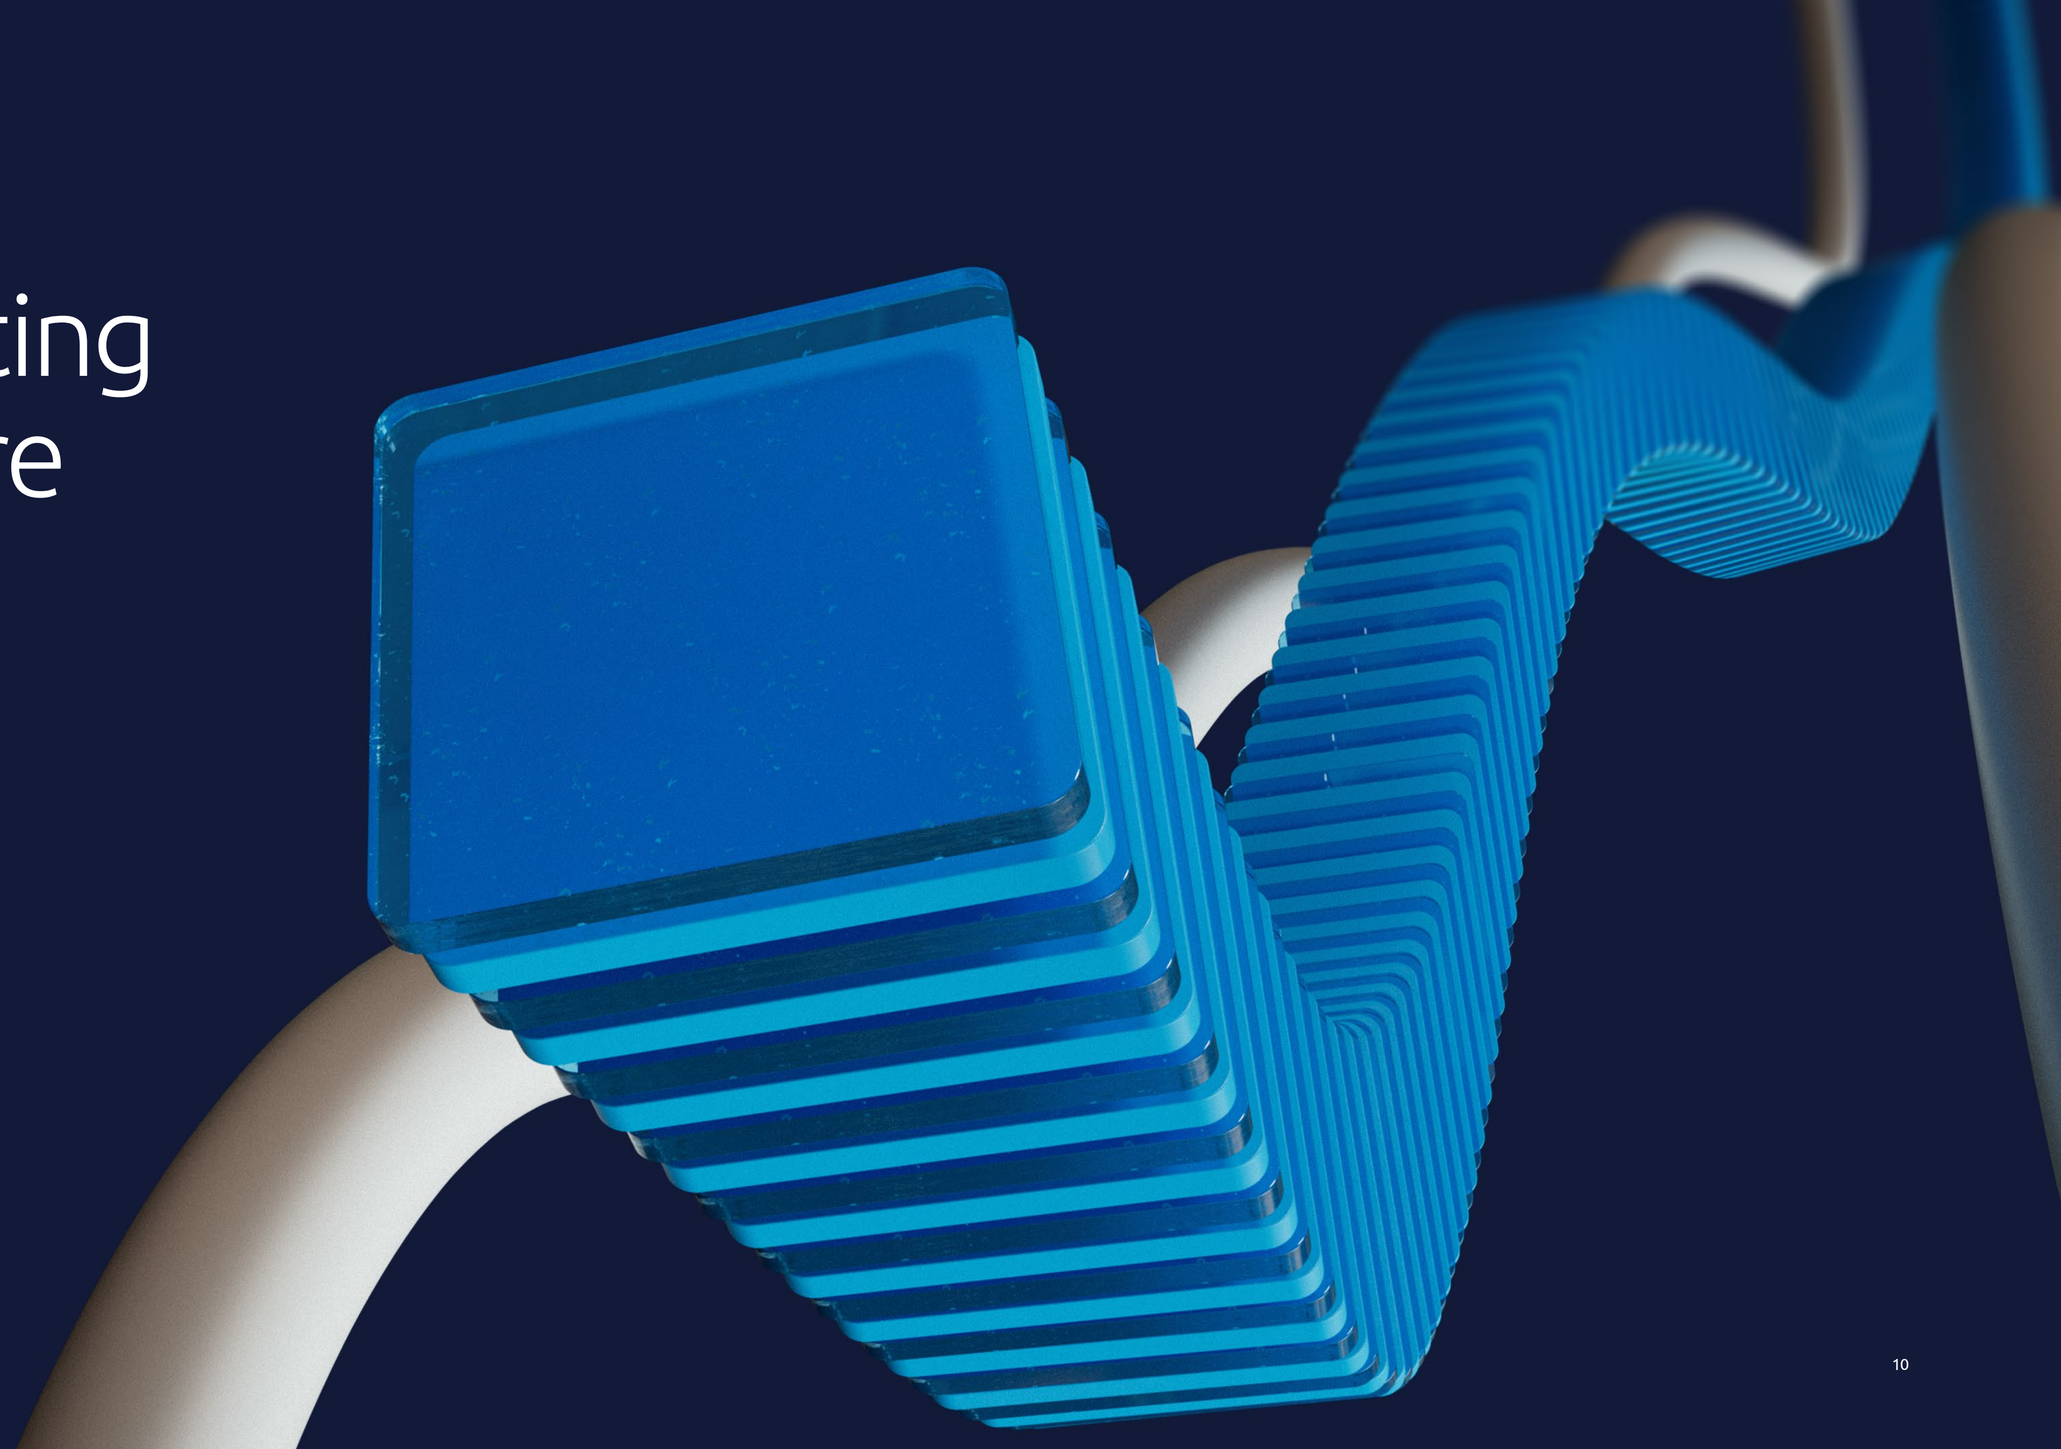

AI is eating software

For more than two decades, software has powered digital transformation. After the era when “software ate the world,” we now enter a new phase where “AI is beginning to eat software itself”. What started as isolated AI tools (code completion, automated testing, prompt-based generation) has evolved into a fundamentally new software development paradigm, where humans and AI continuously conceptualize, design, build and refactor systems together.

AI-native development is no longer experimental. Large enterprises, in particular, are moving first: three-quarters of organizations with more than $20 billion in annual revenue have already piloted or scaled generative AI for software[3] engineering. Enterprises are moving beyond AI-assisted coding to fully autonomous software development ecosystems, deploying self-directed testing frameworks, intelligent code-generation agents, continuous auto-refactoring engines, and agentic build-and-release systems that collaborate, learn, and optimize with human interventions. The impact is transformative: delivery cycles accelerate, technical debt is resolved earlier, and developers shift from manually writing code to expressing intent and orchestrating intelligent pipelines.

In 2026, this shift becomes structural. Developers increasingly describe outcomes in natural language or high-level specifications. AI generates the design, implements it, tests it, secures it, optimizes it, integrates it and continuously refactors it as requirements evolve. A new model emerges where AI-assembled components adapt in near real time and software becomes an evolving service rather than a static asset.

At this point, the traditional concept of an application will begin to fade. Users express goals, and AI agents dynamically assemble, run, and maintain the underlying logic. The visible app layer shrinks; behind it, composable AI services evolve autonomously, self-testing, self-healing, and updating themselves as new conditions arise.

This transition brings enormous opportunity. It also demands a fundamental unlearning of old software habits. AI-generated code is powerful but not infallible, making governance, validation, and architectural oversight more critical than ever. The skills that once differentiated developers (package configuration, front-end coding, manual quality assurance) lose importance. The new currency of expertise becomes systems thinking, AI orchestration, architecture, and the ability to manage complex autonomous toolchains.

Why it matters

This shift is not only technological; it is strategic. As enterprises reach the limits of traditional DevOps, AI-native development introduces a new pathway to speed and agility. Agentic pipelines can generate, test, and refactor software continuously, shrinking quality assurance cycles and enabling near-real-time adaptation. This is a foundational element in building AI-native businesses. When software evolves automatically rather than through manual releases, systems become adaptive, allowing organizations to respond to market changes far faster than static architecture currently allows.

At the same time, AI-generated software opens new options for digital sovereignty. By lowering the cost and effort required to design, maintain, and evolve software, AI-native development reduces the economic barriers that have historically pushed organizations toward large, standardized SaaS platforms. This makes it viable—where strategic control, regulatory constraints, or data locality matter—to replace monolithic SaaS with tailored systems whose codebase, data flows, and evolution remain under direct organizational control and aligned with strategic autonomy goals.

Finally, this shift frees human talent for higher-value work. Automation of routine software development lifecycle tasks redirects engineers' focus toward architecture, product logic, and governance, provided they are able to build on traditional practices and master new AI-driven toolchains.

Together, these forces redefine how software is built, how fast it evolves, and how enterprises differentiate in an AI-driven era.

What to look out for

Enterprises begin deploying **agentic build systems** that generate, test, secure, and refactor code continuously. These systems go far beyond today's code assistants and act as full orchestration engines, translating intent into production-ready code and maintaining it as requirements evolve. Early adopters already use them to reduce technical debt and compress release cycles dramatically.

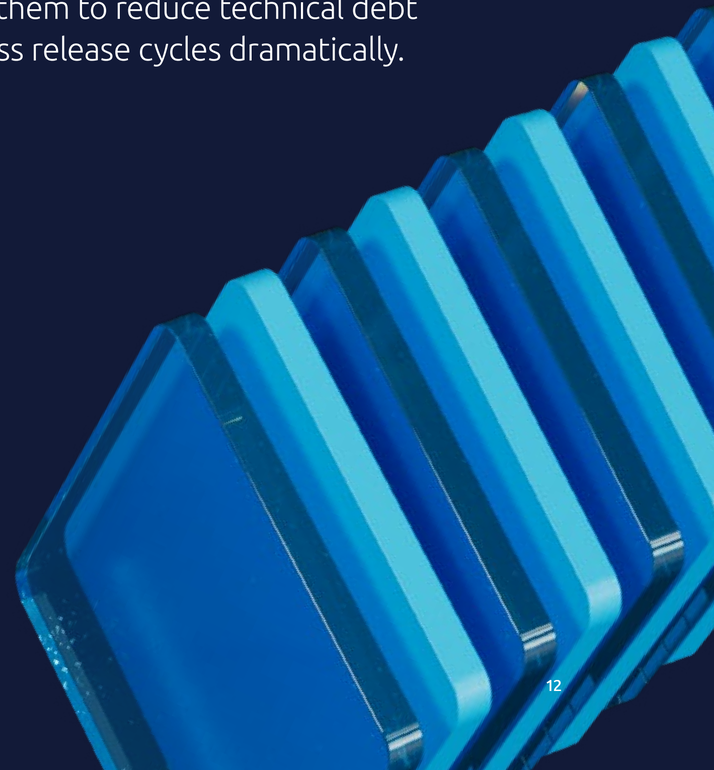

We also see rapid advances in **autonomous quality assurance and reliability pipelines,** where test generation, regression detection, vulnerability scanning, and dependency management are handled end-to-end by AI. As manual QA becomes impractical at the speed AI enables, these reliability pipelines form the backbone of trustworthy AI-generated software.

Another signal of this shift is the rise of **dynamic, composable services** assembled automatically by AI. Instead of treating applications as fixed assets, enterprises move toward adaptive Service-as-Software models where components are assembled, optimized, and updated in near real time. This enables greater differentiation and supports the evolution toward sovereign, custom-built systems.

Finally, **AI-governed development environments** become essential to manage this transition. Evaluation frameworks, behavioral monitoring, lineage tracking, and architectural guardrails mature to ensure autonomous toolchains remain reliable, controllable, and aligned with enterprise requirements.

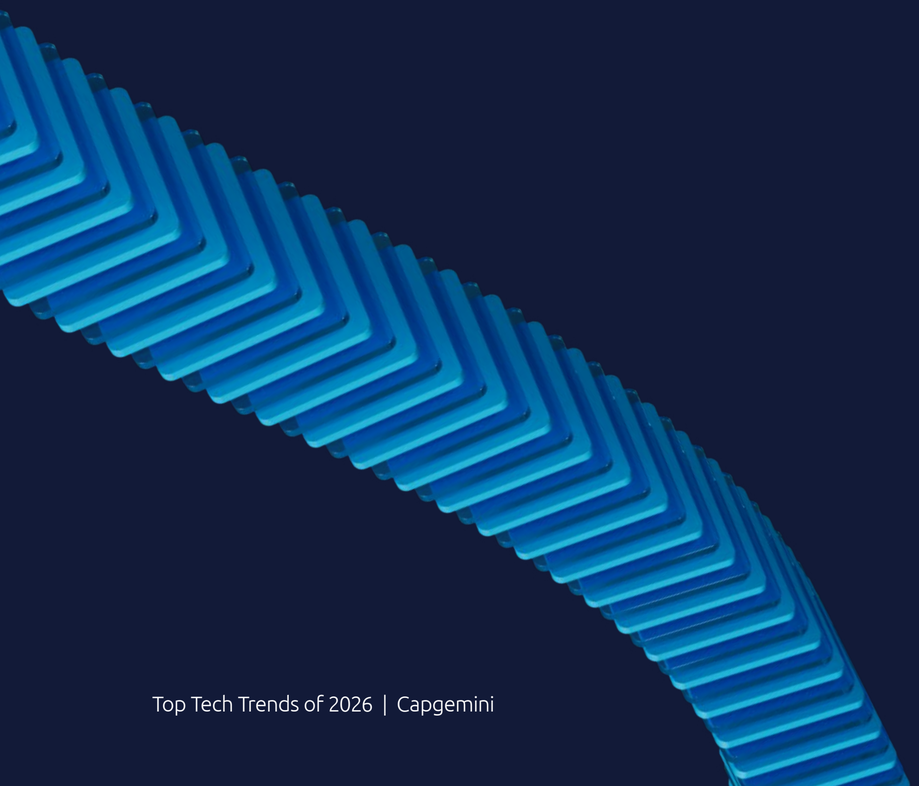

“

The fundamentals of software creation are being rewritten. We are entering an era of AI powered software rebuild, where competitive advantage hinges on deep control over costs, sovereignty, cybersecurity, and data privacy.”

Sudhir Pai

Deputy Group CTO, Capgemini

The rise of intelligent ops

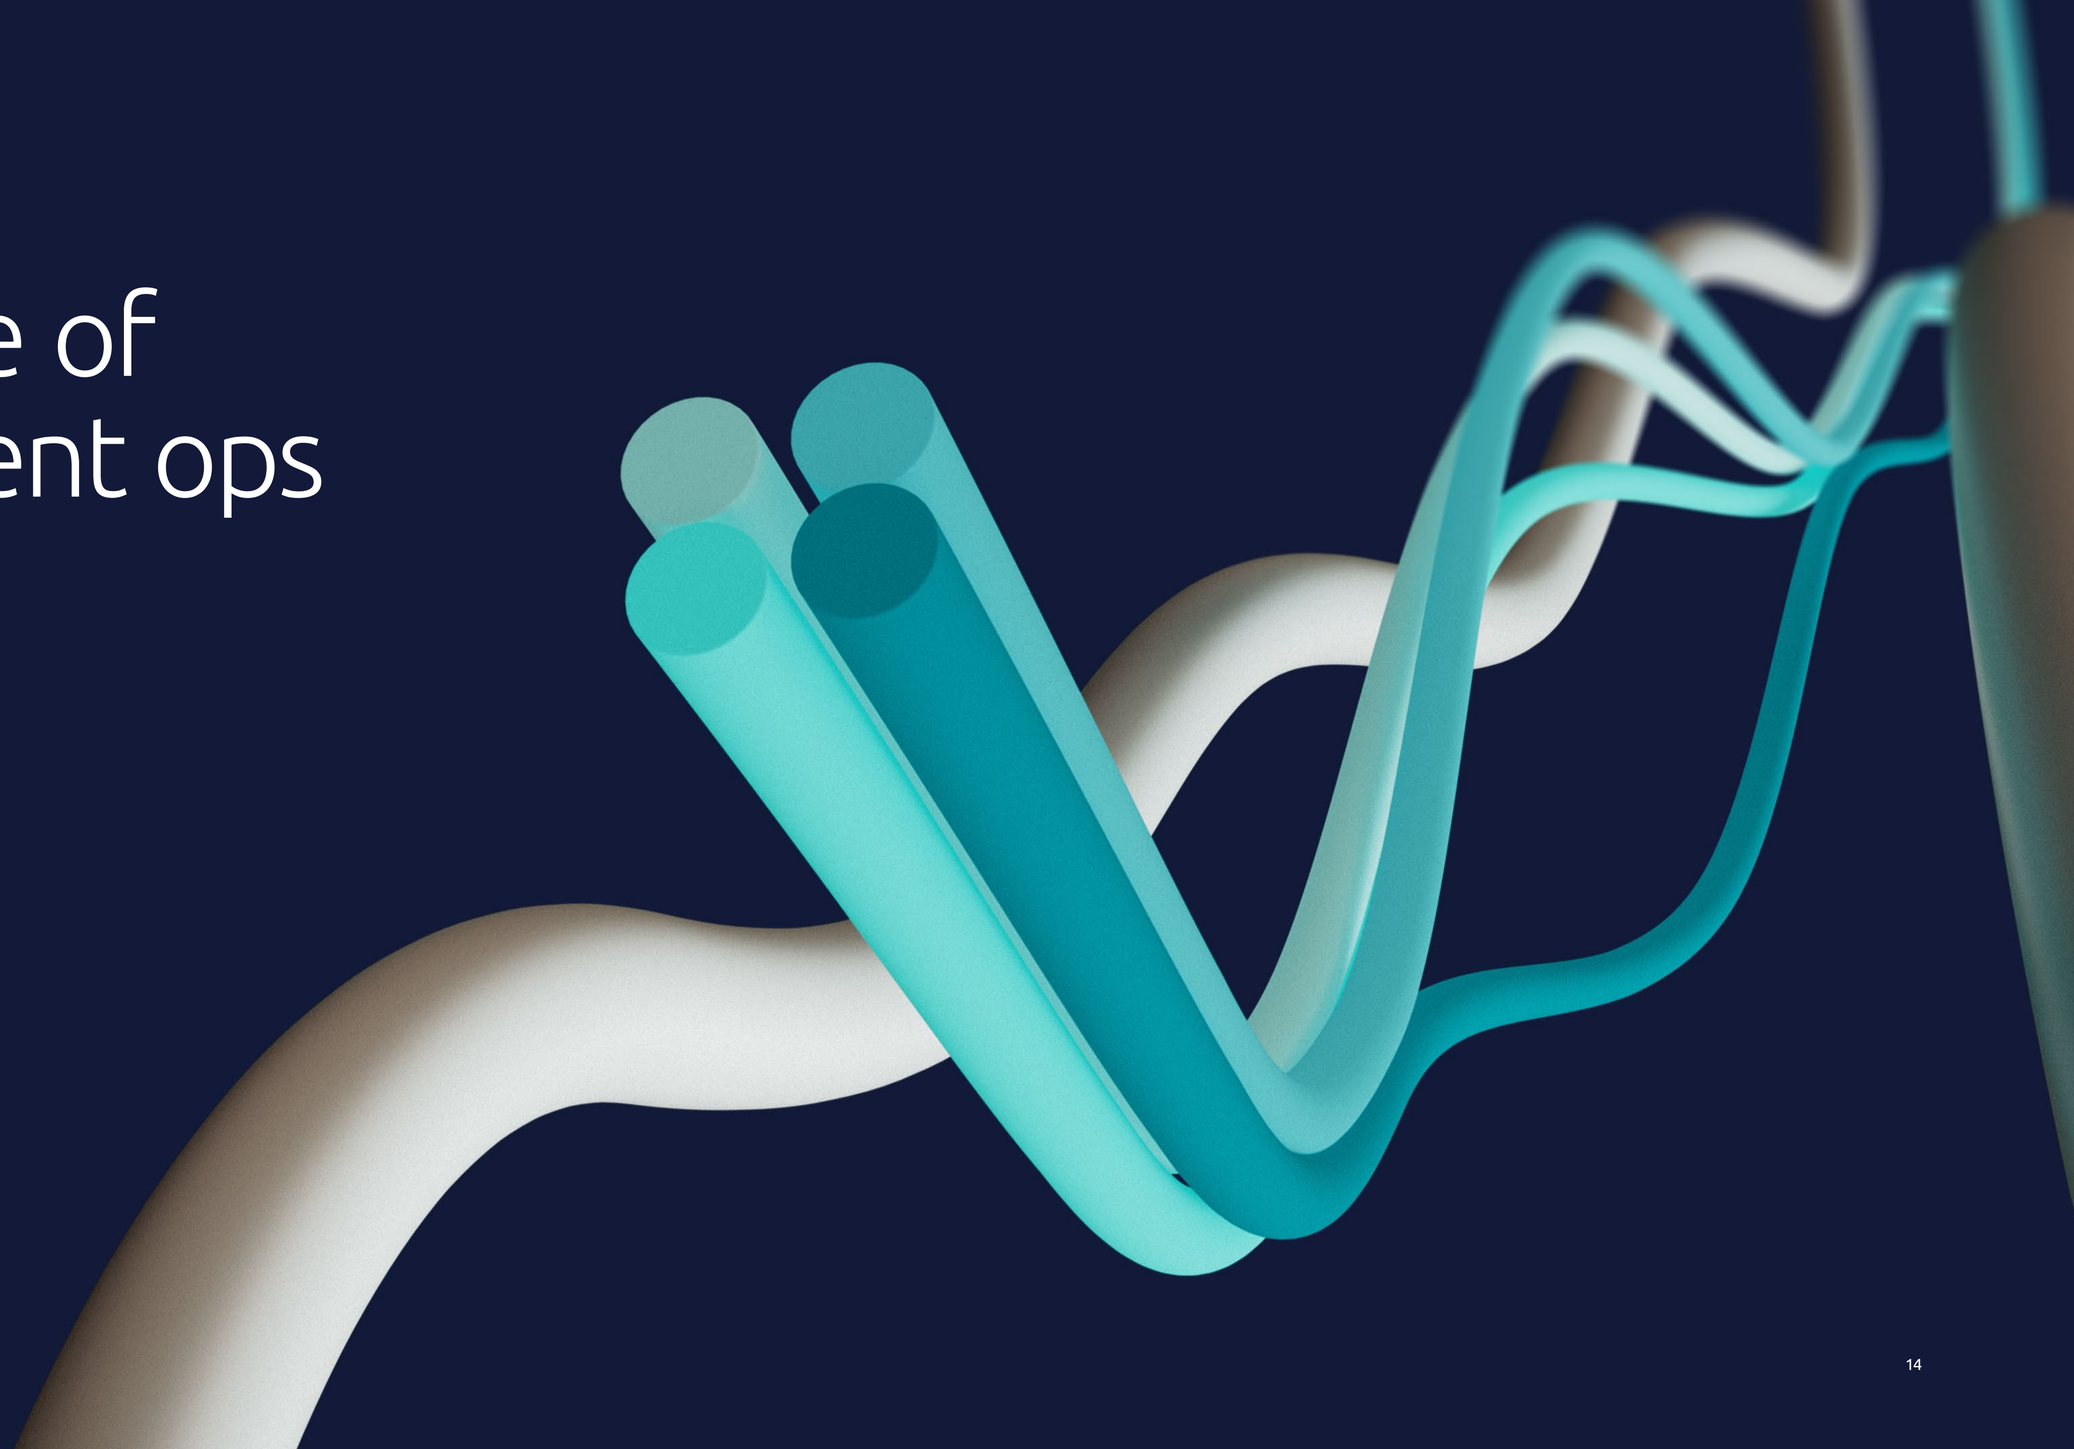

The rise of intelligent ops

After years of enterprise resource planning (ERP) modernization, cloud migration, and process automation, enterprises are entering a new era—one where their operational core becomes intelligent by design. AI is no longer bolted on; it is enabling the complete redesign of business processes, transforming them from linear, predefined sequences into living systems capable of continuous optimization.

In 2025, the use of AI agents in operations more than doubled to 21% (up from 10% in 2024)[4]. The speed of AI adoption, the explosion of real-time data, and the emergence of autonomous agents are driving the rise of intelligent operations. At the center of this shift sits a new kind of enterprise core: dynamic, modular, and powered by data, where platforms, AI agents and humans collaborate in real time to deliver efficiencies and business value while maintaining operations that are agile and resilient. By 2028, Gartner expects close to 40% of enterprise workflows to be automated or augmented by AI agents, underscoring how deeply intelligence is moving into day-to-day operations[5].

Three structural shifts define this transition. First, intelligent processes are being re-imagined to make them AI ready. The AI first processes are hyper automated using a blend of technologies spanning across RPA, Gen AI and agentic AI. Together, these technologies give processes the ability to interpret signals, adapt workflows, and trigger actions in an intelligent way. This direction is already explicit in executive intent: 82% of organizations plan to integrate AI agents by 2027, and 85% of executives expect them to autonomously handle one or more business processes within the next three to five years[6].

Second, organizations are moving beyond isolated process optimization to orchestrating value chains. This holistic approach breaks down silos, eliminates friction at process handoffs, and creates seamless, end-to-end operations that connect functions into a unified value stream.

Third, the role of people evolves toward human–AI co-steering: AI proposes and executes, while humans supervise, validate outcomes, and apply judgment. This partnership blends machine speed with human insight to accelerate execution while safeguarding safety, ethical guardrails and trust.

Why it matters

Operations are the heartbeat of every organization—the fundamental activities that keep a business running. They are data-rich, repetitive, and deeply tied to business outcomes, making them one of the most powerful use cases for AI.

In 2026, leaders face a breaking point: they must move beyond pilots and isolated automation to match the pace of change. Intelligent operations offer a scalable path to transform the core of the organization and unlock greater value. The payoff is significant: continuously adapting operations, reduced friction, and a digital core that evolves with the business rather than slowing it down. Intelligent operations ensure enterprises not only run better, but allow them to continuously reinvent themselves, marking the moment when AI becomes a true co-operator.

In addition to efficiency, this shift is about redefining how businesses create business value. Intelligent operations unlock growth by enabling faster innovation, improving customer experiences, and building resilience into every process. When AI becomes embedded in the operational core, enterprises gain the agility to anticipate change, adapt in real time, and turn disruption into advantage. That's why operations represent one of the most transformative opportunities in 2026 and beyond.

What to look out for

In 2025, agentic AI projects surged by 48%, driven largely by experimentation through pilot initiatives[7]. In 2026, expect more and more organizations to move from pilots to first production levels, with the focus evolving from cost savings to value creation.
One of the clearest indicators is the growing intent to replace legacy workflows outright with AI-enabled ones. **This isn't about layering automation on top of old processes. It's about redesigning operations from the ground up,** with generative AI, RPA, analytics, and AI agents working hand in hand. The goal isn't incremental improvement; it's a bold reset. We see this happening in consumer products, where leading organizations are making bold moves to embed AI into their sales and operations planning: AI predicts and forecasts, then dynamically adjusts manufacturing, distribution, and logistics to boost sales and reduce stock.

A second tell is the rise of **connected operations** that span entire value chains. Projects in which supply chain, finance, procurement, and customer operations operate as one orchestrated system (with AI coordinating handoffs and resolving issues across functions) show that organizations are moving beyond isolated KPIs and beginning to treat processes as integrated flows.

Finally, another clear signal is **human–AI co-steering** in operational command centers becoming more formalized and strategic. Instead of dashboards that only report metrics, teams start to work in environments where AI suggests actions, executes routine tasks, and escalates only when human judgment is required. Pilots already show AI handling month-end close steps and vendor risk checks which frees up humans to focus on exceptions and strategic decisions that drive value.

Together, these projects demonstrate that operations are shifting from fragmented workflows to adaptive, AI-operated systems. The real story beyond cost reduction, is the ability to unlock new value streams and transform operations into a true growth engine.

"Operations are no longer **static systems** optimized once and maintained over time. They are becoming intelligent, adaptive cores—where AI is embedded into processes and **humans co-steer systems that learn and evolve continuously.**"

Manuel Sevilla

CTInO, Capgemini Business Services

Cloud 3.0 – all flavors of cloud

Cloud 3.0 – all flavors of cloud

Cloud is entering its next evolution. After a decade focused on migration, modernization, and cost optimization, it is becoming the execution backbone for AI and AI-assisted applications. This shift is reflected in scale: public cloud spending is expected to almost double from $723 billion in 2025 to $1.47 trillion by 2029, with generative AI projected to account for 10–15% of that spend by the end of the decade[8]. Cloud value is no longer defined by elasticity alone, but by its ability to sustain AI-intensive, always-on workloads.

Yet AI cannot scale on classical public cloud architectures alone. The need to fine-tune models on proprietary data, control data sensitivity, meet regulatory constraints, and deliver low-latency inference is pushing organizations toward hybrid, private, multi-cloud, and sovereign cloud architectures, not as exceptions but as the new norm. Agentic systems, in particular, require scalable, distributed, low-latency infrastructure where cloud and edge operate as a single intelligent fabric.

This shift is reinforced by two structural pressures. First, large-scale cloud outages have exposed the risk of single-provider dependence, prompting many enterprises to design for portability and resilience.

Second, geopolitical tensions and sovereignty requirements are reshaping cloud choices as organizations seek greater control over where data resides and how workloads are executed.

By 2026, hybrid platforms will become mainstream. AI workloads will flow seamlessly between edge, private cloud, and public cloud environments. Organizations will redesign their architectures for resilience, interoperability, and strategic autonomy. Cloud 3.0 marks the moment when cloud stops being only an IT platform and becomes the distributed execution layer that makes AI (especially agentic AI) possible at scale.

Why it matters

Cloud 3.0 increases resilience, but also complexity. As AI adoption accelerates, enterprises must operate confidently across diverse environments, providers, and jurisdictions. The differentiator is no longer choosing a single platform but mastering interoperability, portability, and intelligent workload placement.

This shift also raises expectations for cloud providers. They will need to remain agile regarding multi-vendor flexibility, sovereign deployment options, cross-cloud compatibility, and seamless edge-to-cloud integration, all while supporting the performance demands of generative AI and agentic workloads. Providers able to deliver openness and resilience will shape the next phase of the market.

For enterprises, Cloud 3.0 changes the operating model from managing a centralized cloud platform to orchestrating a continuously evolving, distributed computing fabric. New skills, governance, and product thinking are required to orchestrate distributed infrastructures, ensure continuity, and deploy AI workloads safely across clouds and edges. Those who adapt turn cloud from a cost center into a strategic enabler of speed, autonomy, and competitive advantage.

Cloud 3.0 is therefore not just another stage of cloud adoption: it becomes the foundation for running AI at scale, building operational resilience, and enabling the next decade of intelligent enterprise architectures.

What to look out for

Several concrete developments signal that Cloud 3.0 is taking hold in 2026. One early indicator is the emergence of **AI-optimized hybrid architectures,** where enterprises deliberately split workloads across private, public, and edge environments. Projects where fine-tuning runs in controlled private clouds while inference executes at the edge for latency-sensitive applications show that AI workloads are already moving fluidly across environments instead of defaulting to a single hyperscaler—a clear sign that Cloud 3.0 is underway.

A second tell is the adoption of **multi-cloud portability frameworks.** Enterprises now run critical applications on two or more cloud providers, not for cost reasons, but for resilience and sovereignty. Containerized AI services, cloud-agnostic data layers, and replicated inference stacks deployed simultaneously across providers become visible proof points, emerging in response to recent outages and geopolitical risk.

Another signal is the **deployment of sovereign cloud regions or sovereign overlays** for sensitive workloads. Financial services, public sector, and healthcare organizations are leading pilots that combine hyperscaler scale with strict locality controls, audited administration, and independent cryptographic governance.

Together, these projects show the shift from cloud as a destination to cloud as a distributed, intelligent substrate for running AI at scale. These are the tells that Cloud 3.0 has entered the operational mainstream.

“Cloud is entering a new phase, driven by the need to provide the intelligent fabric required to run AI workloads and orchestrate agentic systems. This future cloud is inherently multi-faceted, spanning multi-cloud, hybrid, and edge environments operating as a single continuum.”

Eric Fradet

CTO and Head of Industrialization of Cloud Infrastructure Services, Capgemini

The borderless paradox of technological sovereignty

The borderless paradox of technological sovereignty

Escalating tensions around critical supply chains, renewed scrutiny of data and cloud dependencies, and growing concern over AI infrastructure control have pushed sovereignty back to the forefront—at levels not seen since the COVID-19 crisis. Enterprises like governments are recognizing how deeply modern economies rely on global technology ecosystems they do not fully control.

At the same time, the definition of "sovereignty" is far from settled. Does it mean localizing data, operating domestic cloud infrastructure, building indigenous chips, mandating open-source alternatives, or reducing dependence on foreign AI models? And to what extent is any of this achievable? In a world built on inherently borderless technologies, complete technological autonomy may appear to be an illusion. Modern digital systems depend on globally distributed supply chains—from semiconductors and hyperscaler cloud platforms to open-source frameworks and frontier AI models. No single organization can realistically detach from this web without losing innovation, scale, or competitiveness.

This leads to a new reality: sovereignty is no longer defined by isolation, but by resilient interdependence.

For most business leaders, the conversation has shifted toward continuity and control: ensuring that critical operations cannot be disrupted by sanctions, outages, geopolitical shocks, or decisions made by non-sovereign actors. Sovereignty becomes a question of resilience: the ability to guarantee secure access to data, infrastructure, and talent regardless of external events.

In response, governments and enterprises are accelerating investment not in full decoupling, but in credible alternatives that reduce dependency risks. Regional cloud providers, sovereign cloud offerings, domestic chip initiatives, and open-AI ecosystems are gaining momentum. All major hyperscalers have now launched or announced sovereign and regulated cloud offerings for 2026, reflecting a structural shift toward sovereignty-by-design rather than withdrawal from global platforms. These options do not replace global platforms, but they strengthen negotiation power, increase strategic flexibility, and reduce operational fragility.

This mindset results in new architectural principles. Organizations diversify suppliers, adopt multi-cloud and multi-vendor strategies, maintain data portability, and redesign systems so they can withstand regulatory shifts or supply-chain disruptions without interrupting operations. Sovereignty becomes a design choice, embedded into cloud strategy, data management, and AI deployment. The outcome is a new global landscape where infrastructures remain open and interconnected but no longer fragile.

Why it matters

In 2026, the world's major regions will intensify their race for control over the critical layers of the digital value chain: semiconductors, data infrastructure, cloud, connectivity, and AI models. At the same time, hyperscalers and large cloud providers are launching sovereign and region-specific offerings, reshaping how enterprises think about trust, compliance, and operational continuity.

For companies, this shift has profound implications. Dependency without safeguards is no longer acceptable. Technology strategies must now account for geopolitical volatility, supply-chain risk, and regulatory divergence. Architectures must be designed for resilience—ensuring portability across providers, continuity during outages, and flexibility to adapt to new localization rules.

Ultimately, technology sovereignty is no longer about retreating behind borders. It is about governing interdependence: building systems that remain globally connected, but controllable; open, but not exposed; scalable, but resilient. This new approach will define how organizations build trust, manage risk, and ensure long-term competitiveness in the decade ahead.

What to look out for

The strongest signals that technology sovereignty is becoming real in 2026 come from concrete projects appearing simultaneously across cloud, AI, compute, and supply chains. One clear indicator is the rapid adoption of **sovereign cloud and data-control architectures.** Enterprises are moving sensitive workloads to region-governed cloud zones or sovereign overlays, where administration, encryption, and data locality remain under jurisdictional control. When regulated sectors begin running AI inference or critical data pipelines inside sovereign environments rather than defaulting to global hyperscalers, sovereignty is already materializing.

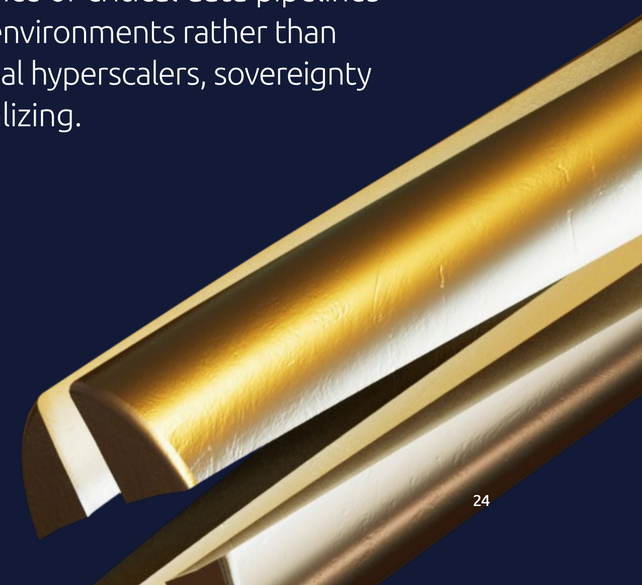

A second tell is the rise of **sovereign AI ecosystems.** Governments and industries are developing locally governed foundational models, national training datasets, and sector-specific AI baselines. These initiatives ensure that high-risk workloads—public services, healthcare, finance, and defense—do not depend solely on foreign black-box models. The moment enterprises start adopting "trusted AI stacks" with clear provenance, auditability, and jurisdictional guarantees, sovereignty moves from policy aspiration to operational practice.

A third signal comes from **domestic compute and resilient supply-chain programs.** Nations and industries are investing in chip fabrication, advanced packaging capacity, sovereign GPU clusters, and diversification of hardware suppliers. Enterprises increasingly design for multi-cloud continuity and hardware portability to protect against sanctions, outages, or geopolitical shock. When organizations begin allocating workloads or procurement choices based on sovereignty risk, the shift is unmistakable.

Together, these developments show sovereignty expanding across technology domains. A move from dependency to managed interdependence will become a defining strategic posture for enterprises in 2026.

"**Technology sovereignty is no longer a policy abstraction.** It has become a strategic concern not only for nations, but for organizations as well. The conversation is shifting from the illusion of isolation **to the practical challenge of managing interdependence and achieving strategic autonomy.**"

Nicolas Gaudillière

CTO, Capgemini Invent France

Emerging signals to watch by 2030 and beyond

Some of the most consequential technological shifts of the next decade will unfold far from the public eye. They will not originate from new digital interfaces or consumer platforms, but from advances in our ability to compute, model, and engineer the physical world at the atomic and molecular scale. Many of the hardest scientific and industrial challenges—from corrosion and material degradation to catalytic reactions, battery stability, and semiconductor reliability—are governed by quantum-mechanical effects that classical methods can only approximate.

A new generation of capabilities is now changing this equation. High-performance computing, AI and foundation models for science, and increasingly automated laboratories are already enabling far more precise exploration of complex physical systems. Progress in quantum hardware and quantum-inspired algorithms is expected to further extend these capabilities over time. Together, these technologies are shifting material innovation away from empirical trial-and-error toward modelling and design grounded in first principles.

These advances matter because they reshape the foundations of energy, infrastructure, manufacturing, sustainability, and health. The impact will not come from the disruption of entire industries overnight, but from a deeper transformation of the materials, components, and physical systems on which those industries depend. Taken together, three emerging signals illustrate how this new wave of materials innovation is unfolding: **materials are increasingly defined by their quantum behavior;** their **design processes now span the full lifecycle,** from manufacturability and performance to degradation and recyclability; and they are progressively **engineered for specific functions** rather than discovered through empirical trial and error. These scientific frontiers point to how the next era of innovation may take shape.

Materials defined by quantum behavior

The next wave of materials innovation will be driven by a deeper understanding of matter at the atomic scale. As technologies push materials to their physical limits, performance, durability, and reliability are increasingly determined by quantum-scale effects that classical engineering approaches struggle to capture. Many long-standing industrial challenges—such as corrosion, degradation, and material fatigue under real operating conditions—originate at this level.

Advances in computation are changing this dynamic. The combination of high-performance computing, AI-driven modelling, and emerging quantum techniques is allowing scientists and engineers to observe, simulate, and predict material behavior with far greater accuracy. This makes it possible to design materials based not only on simplified assumptions, but on how they behave in complex, real-world environments.

The impact is twofold. On one hand, this opens the door to new classes of materials with breakthrough properties, from advanced batteries and catalysts to novel electronic and superconducting materials. On the other, it enables these innovations to be engineered for stability, manufacturability, and long-term use. Together, these capabilities turn quantum-scale understanding into materials that can be deployed at an industrial scale.

Next-generation biodegradability

A second emerging signal is the growing ability to understand and model how materials degrade at the atomic and molecular level. Biodegradability has traditionally been treated as an approximate property, assessed through empirical testing and often achieved at the expense of durability or consistency. As a result, material degradation has remained difficult to predict, control, or design with precision.

Advances in computation are beginning to change this. High-performance computing, AI-driven models, and, over time, quantum computing make it possible to simulate the chemical and physical processes that govern how materials interact with their environment and break down over time. This deeper understanding allows engineers to move beyond trial-and-error approaches and start designing degradation pathways intentionally, based on how materials behave at the atomic scale.

Over the long term, this capability could fundamentally reshape sustainable materials design. Instead of retrofitting biodegradability after performance has been optimized, future materials can be engineered from the outset to balance performance, durability, and controlled end-of-life behavior. This shift has far-reaching implications for industries facing increasing sustainability and regulatory pressures, from packaging and textiles to electronics and advanced materials.

Synthetic material science

The third frontier is the shift from discovering materials to deliberately designing them. Instead of searching for useful materials in nature or through trial and error, scientists are increasingly able to define the properties they want—such as strength, conductivity, or resistance to heat—and design materials to match those requirements. Materials begin to resemble engineered systems, shaped by intent rather than chance.

This shift is enabled by advances in computation and automation. AI-driven models, combined with increasingly automated laboratories, allow researchers to explore vast numbers of possible material combinations and work backward from desired outcomes. What once required years of experimentation can increasingly be achieved through continuous, data-driven design loops that rapidly test, refine, and validate new material candidates.

Over time, these capabilities are becoming more accessible. Cloud-based tools, shared computational infrastructure, and more autonomous design systems are lowering barriers to entry, extending advanced materials engineering beyond a small number of elite laboratories. The result is a broad expansion of material innovation, with implications across mobility, energy, infrastructure, electronics, and healthcare.

What does this mean?

Taken together, these signals show that the next phase of technological progress will be driven by the convergence of advanced computation, chemistry, and biology. AI, high-performance computing, and laboratory automation are already reshaping how matter can be explored and engineered, while quantum computing is expected to further extend these capabilities as it matures. In parallel, advances in molecular engineering and synthetic biology are translating digital insight into physical materials with properties designed for specific functions and constraints.

This is where the quiet shifts of the coming decade will originate: not at the level of interfaces or platforms, but in the underlying materials and physical systems on which industries depend. Organizations that recognize this convergence will be better positioned to anticipate long-term technological change, manage emerging resources and sovereignty constraints, and build durable advantages in a world where control over matter itself is becoming a strategic capability.

Concluding remarks

In 2026, the technology landscape will be impacted by the maturation of AI and its integration across software, the evolution of cloud into multi-cloud ecosystems, and the rise of intelligent operations. These trends, driven by AI's pervasive influence, promise to accelerate hyper automation, redefine business models, and enable adaptive enterprises. Alongside emerging priorities such as agentic systems and the delicate balance of tech sovereignty, these developments underscore a shared trajectory toward interconnected and resilient technology frameworks. It's clear that transformative change occurs when multiple technologies converge seamlessly to address complex challenges.

Embracing these trends will be essential for extracting value, sustaining growth and securing a competitive edge in an increasingly fluid market. Organizations that harness this borderless, intelligent future will not only lead in 2026 but also shape the foundation for the decade ahead.

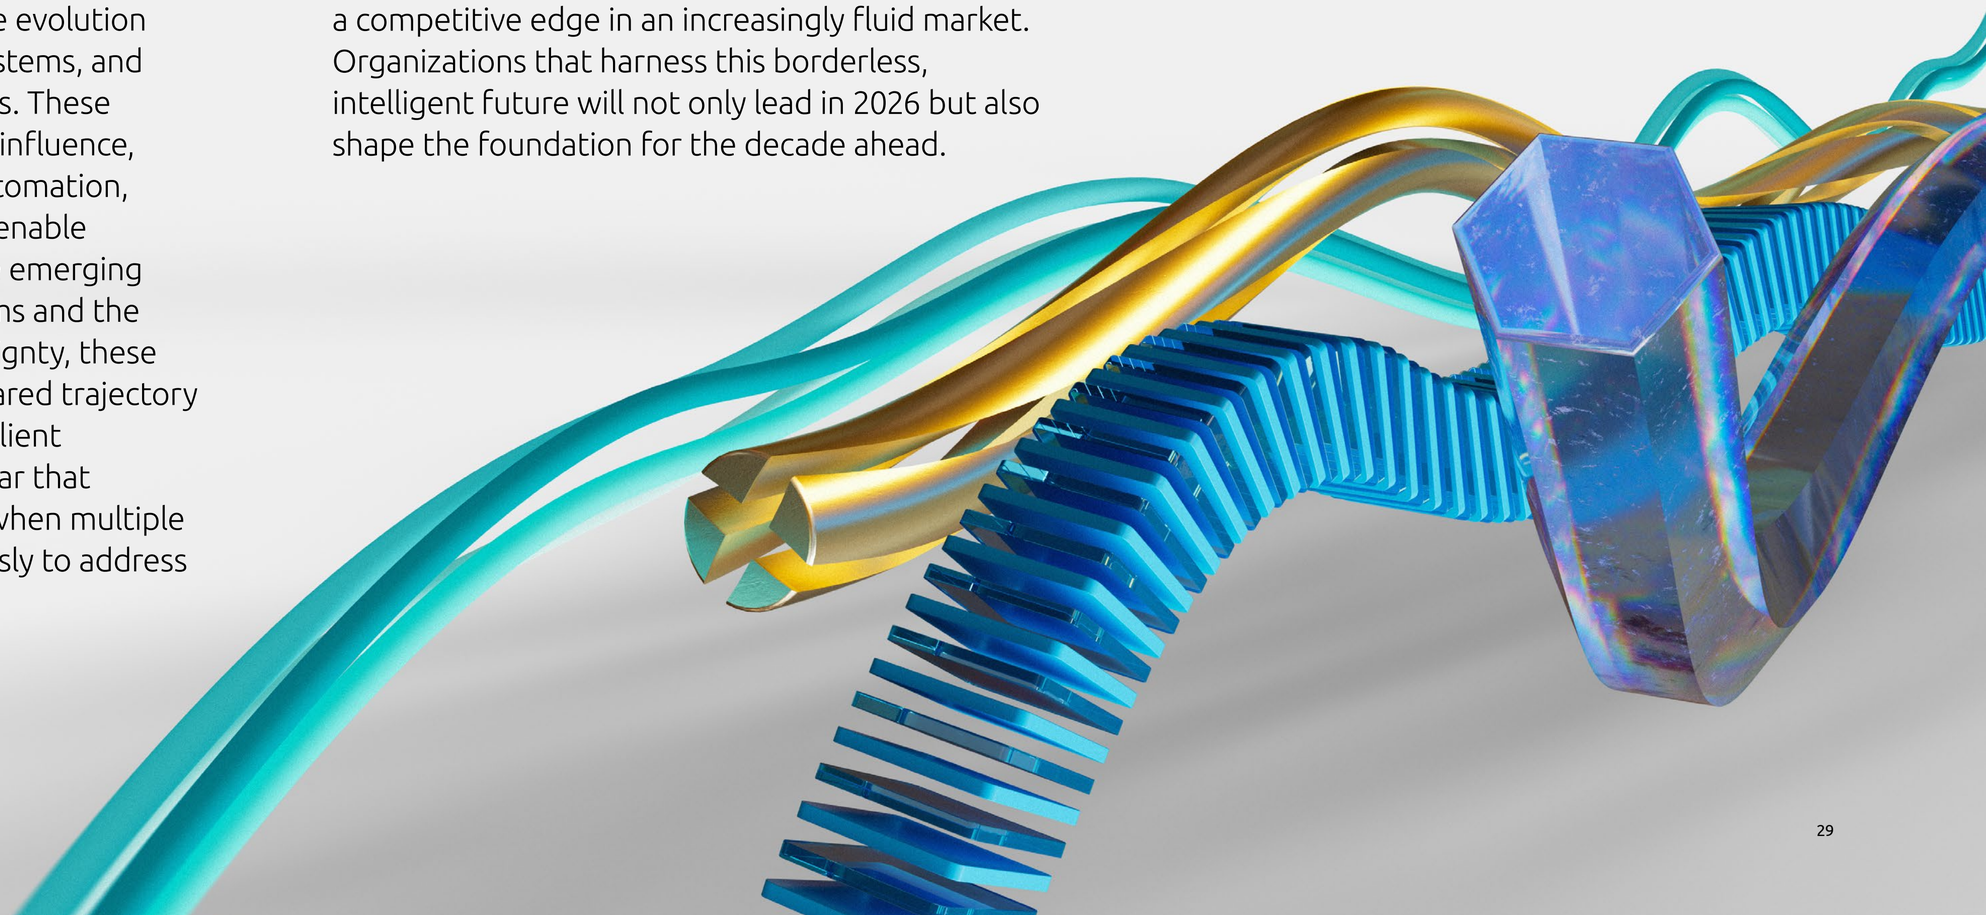

Leading the way in innovation

The pace of innovation is accelerating and the number of emerging technologies hitting the market continues to grow. Their impact is strong and felt globally across industries, in the public and private sectors.

Through TechnoVision, Capgemini's ongoing emerging technology program, our global network of research, technology and innovation teams are at the forefront of innovation: taking the pulse, scouting ideas and identifying which technology trends will be driving transformative change now and in the future.

Our goal is to help our clients assess these technologies and develop sound strategies to make the best investments. We share our vision and provide our clients with the right information to make the right decisions for their business, their people and our planet.

Whether your objective is to better understand the technology trends that will shape the year ahead or you are looking to deep dive into the assessment and planning of a digital transformation, our TechnoVision program is a beacon in an ever-evolving technology landscape and it is an open door into our robust innovation ecosystem.

Our expertise in emerging technology, our academic partnerships, our focus on applied innovation, our design and engineering capabilities, our investment in ventures and collaboration with technology and industry partners have been designed with one thing in mind—to help you unlock the potential of innovation and bring your ideas to life. As business leaders and decision-makers, one of the most important decisions you can make is to stay ahead of the curve and anticipate what will matter tomorrow. Our innovation and technology teams are here to help.

Capgemini's TechnoVision guide, designed to help organizations assess their technology environments, will be published in February 2026.

References

1 Capgemini Research Institute – Turbocharging Software with Gen AI (2024)

2 Capgemini Research Institute – Turbocharging Software with Gen AI (2024)

3 Capgemini Research Institute – Turbocharging Software with Gen AI (2024)

4 Capgemini Research Institute – AI in Action: How Gen AI and Agentic AI Redefine Business Operation (2025)

5 Gartner – Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025 (2005)

6 Capgemini Research Institute – Harnessing the Value of Generative AI: 2nd Edition (2024)

7 Capgemini Research Institute – Harnessing the Value of Generative AI: 2nd Edition (2024)

8 Gartner – Gartner Forecasts Worldwide Public Cloud End-User Spending to Total $723 Billion in 2025 (2024)

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2024 global revenues of €22.1 billion.

Make it real | **www.capgemini.com**