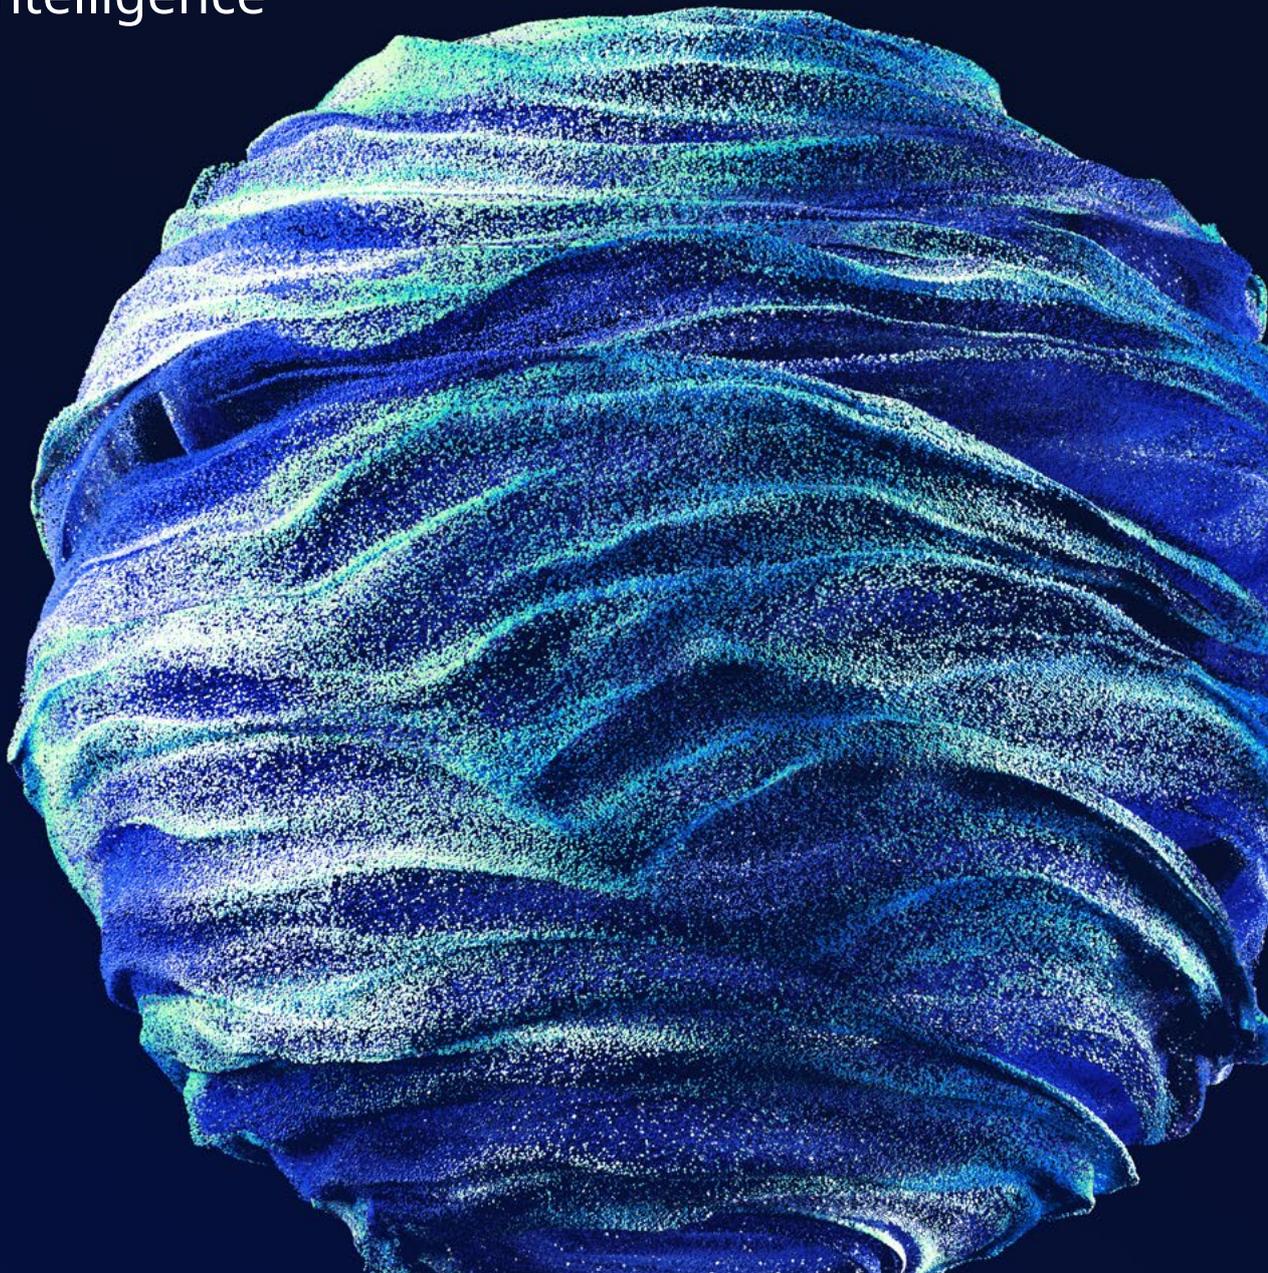


March 2026

NTT DATA Technology Foresight 2026

Sustaining growth in the era
of mass intelligence



Contents

03 Foreword

04 Technology Foresight: A lens into the future

06 Macrotrends overview: The architecture of the future

08 Human-orchestrated autonomy

18 Embodied agency and emotions

28 Intelligence we trust

38 Informed infrastructure

48 Sovereign silicon ecosystems

58 From illusory efficiency to sufficiency

68 Technology in context: Foresight across industries

87 From insight to action

89 About the research

90 Contributors

90 About NTT DATA

Foreword

In the era of mass intelligence, AI has become a capability in the hands of billions, supporting the way people process information, collaborate and solve problems.

Human creativity is now amplified by AI. We see human-orchestrated autonomy transforming organizations: Machines handle tasks at scale while people set direction and purpose. Technology is no longer a cold utility — it is becoming more empathetic. Through embodied agency and emotions, digital agents are beginning to understand human context from nonverbal cues, making interactions more natural and human-centric.

The future we envision is one of intelligence we trust. As AI weaves into decisions from the boardroom to the customer interface, trustworthiness becomes essential. We build trustworthy intelligence by designing algorithms and data ecosystems with ethical responsibility and transparency at their core.

At the same time, our world is supported by informed infrastructure — a smarter, distributed foundation where cloud is only one piece. Urgent decisions are increasingly taken at the edge, demanding workloads run on specialized computing, and sovereignty defines where intelligence lives. Organizations need to build infrastructure that continuously shapes itself around purpose and resilience.

The rise of mass intelligence also demands greater control over the critical elements of the technology supply chain. Sovereign silicon ecosystems are emerging as businesses and nations seek secure, specialized foundations that ensure innovation remains aligned with core values. And as we push the envelope of productivity, the goal is no longer hyperefficiency at any cost, but sufficiency. By shifting from illusory efficiency to sufficiency, leaders prioritize lasting prosperity over short-term gains.

At NTT DATA, our vision is bold and aspirational, yet also pragmatic and grounded in reality. With our own foresight-based AI assistant enhancing how we advise clients, we help organizations turn these forces to their advantage, so collaborative intelligence — the fusion of human intent and machine speed — drives positive change. Those who navigate this era with trust, clarity and purpose will not only adapt to the future but also shape a smarter and more sustainable world.

Technology Foresight: A lens into the future

Purpose and methodology

The purpose of Technology Foresight is to provide strategic clarity amid accelerating technological and societal transformation.

It does not aim to predict a single future but to illuminate multiple plausible futures shaped by deep structural change. Through the continuous exploration of emerging signals, long-term trends and cross-domain dynamics, we help organizations anticipate nonlinear change and understand how global forces may influence markets, technologies and human behavior.

Our methodology combines environmental scanning, expert insight and systemic analysis, integrating social, economic, environmental, geopolitical and technological perspectives. A structured process, accelerated and enhanced by AI, enables leaders to identify early signals, assess implications, and align innovation and investment decisions with long-term societal purpose, ensuring that technological progress supports human and planetary resilience.

These insights have converged into the six technology macrotrends that frame this report. These macrotrends are not isolated developments. Rather, they represent interconnected directions of transformation that influence how intelligence is created, governed, and embedded in and scaled across societies and industries.





Shaping innovation and strategy

Technology Foresight acts as a strategic bridge between insight and action.

By connecting global shifts to emerging technologies and behaviors, it enables organizations to anticipate areas of opportunity and design innovations that meet future needs, not just current demand. And it informs innovation by revealing untapped areas for new offerings, capabilities and business models aligned with likely future scenarios.

Strategically, the insights presented in this report help organizations strengthen resilience by guiding portfolio choices, capability development and ecosystem partnerships. They enable leaders to prioritize investments and determine where to experiment — especially when technological readiness outpaces organizational, regulatory or societal preparedness.

In essence, Technology Foresight enables organizations to navigate uncertainty with confidence, intent and responsibility, ensuring that innovation remains purposeful, future-fit and aligned with societal transformation.

Macrotrends overview: The architecture of the future

We have entered the era of mass intelligence, where every system, process and product is increasingly capable of learning, adapting and acting autonomously, and where AI is now accessible to billions.

NTT DATA Technology Foresight 2026 explores how this transformation can serve purposeful growth through systems that are intelligent yet transparent, emotional yet ethical, and sovereign yet connected.

The six macrotrends in this report represent the architecture of that future: a shift from automation and assistance to autonomy and embodiment through emotionally intelligent and ethically governed systems, supported by informed infrastructure, sovereign hardware ecosystems and a renewed commitment to sustainability through sufficiency.

Together, these dynamics signal the rise of interconnected, adaptive intelligence that will change the way value is created across society.

“

Mass intelligence amplifies human purpose when guided by human values.”

Architecture of the future

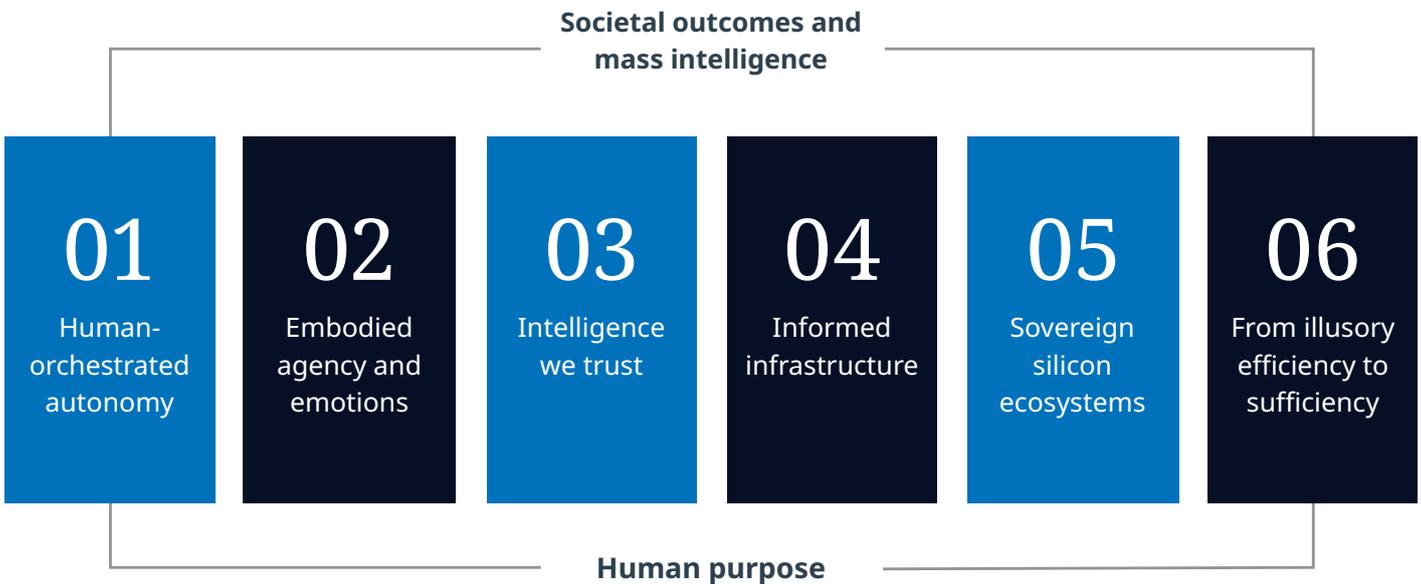


Figure 1: Architecture of the future

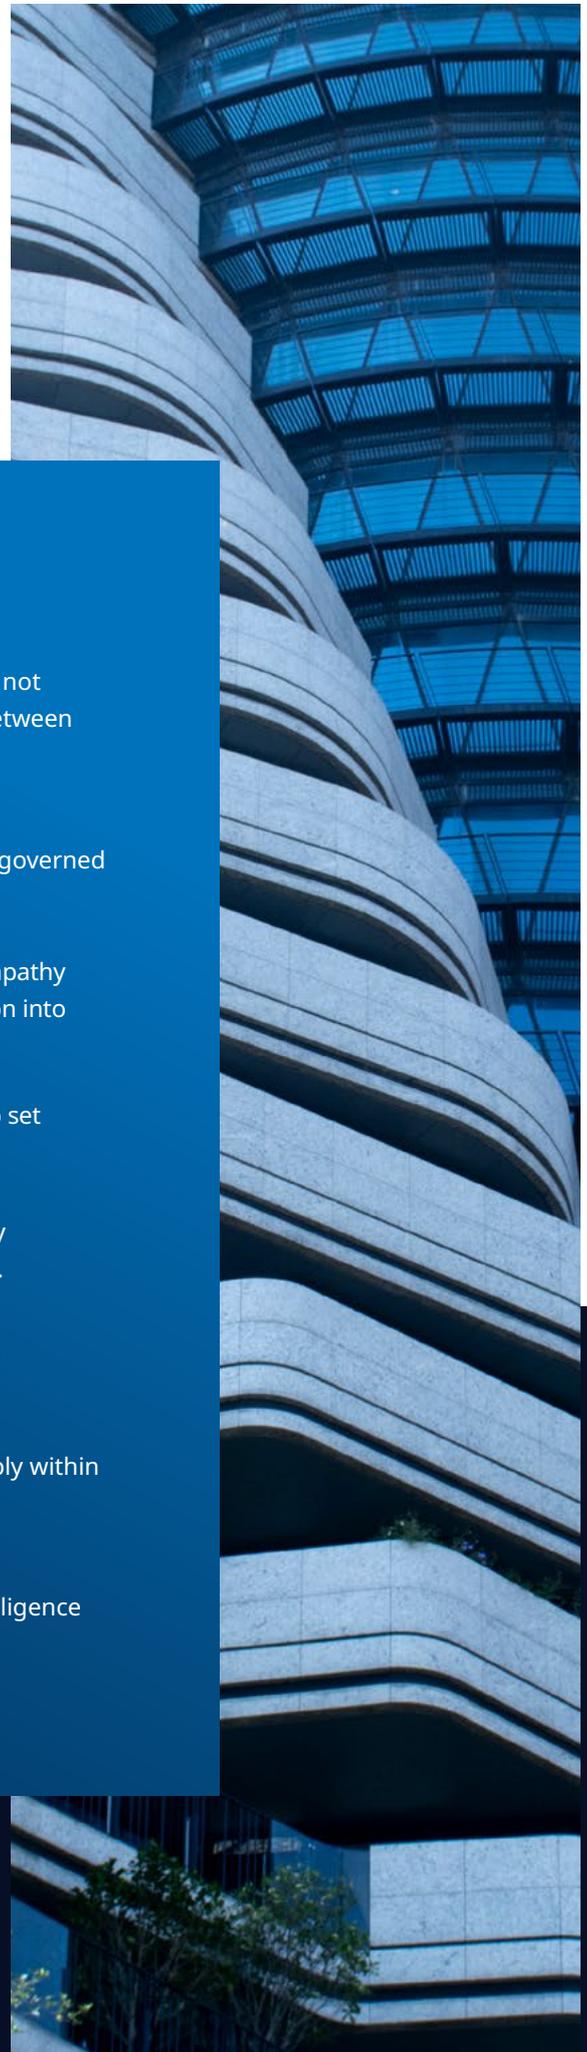
While presented as six distinct macro trends, these forces operate as a single, convergent transformation. Each trend — autonomy, emotion, trust, infrastructure, sovereignty and sufficiency — reinforces the others, forming a coherent foundation for the era of mass intelligence.

Why humans matter in the era of mass intelligence

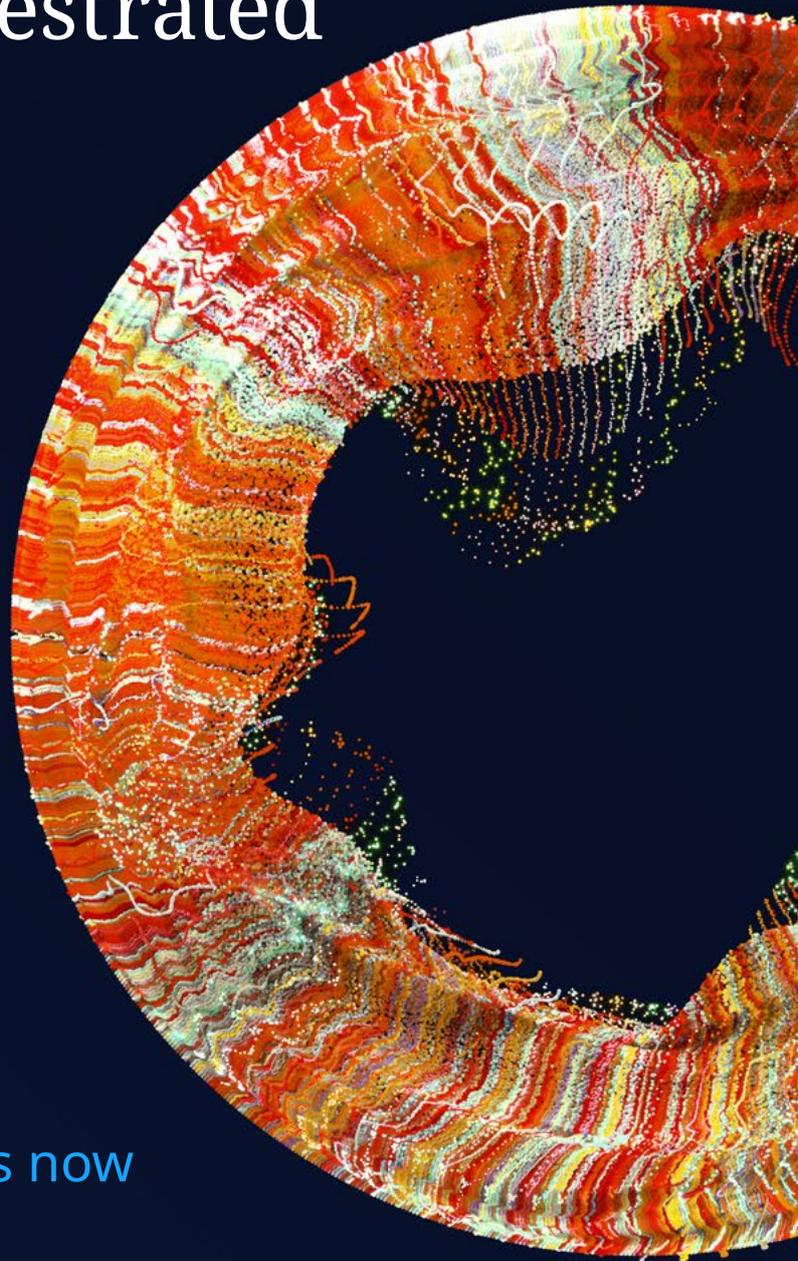
Across all six macro trends, the era of mass intelligence is defined not by machines acting independently but by deeper collaboration between humans and intelligent systems.

- 01 Autonomy** scales human intent through accountable, governed decision-making.
- 02 Embodied agency** enhances interaction by bringing empathy and emotional awareness into digital systems, and action into the physical world.
- 03 Trustworthy intelligence** relies on human judgment to set boundaries of safety, fairness and accountability.
- 04 Informed infrastructure** supports human priorities by enabling resilient, sovereign and adaptive foundations.
- 05 Sovereign silicon** secures the computing capacity that underpins human-directed innovation.
- 06 Sufficiency** ensures that intelligence evolves responsibly within ecological and societal limits.

Together, these macro trends reinforce a core principle: Mass intelligence amplifies human purpose when guided by human values.



1 Human-orchestrated autonomy



What it is and why it matters now

Autonomous systems are a core building block of mass intelligence

Autonomy has reached an inflection point. What began as task automation is rapidly evolving into purpose-led, adaptive intelligence that's capable of taking action across business functions, physical systems and social infrastructure. Yet this acceleration comes at a moment when public trust, regulatory pressure and ethical expectations are rising sharply.

This makes human-intent alignment the new strategic frontier. Organizations can no longer rely on opaque algorithmic behavior: They need autonomous systems that are transparent, explainable, auditable and continuously governed.

As enterprises move from pilot projects to scaled autonomous operations — whether in robotics, decision automation or AI-native agents — the critical challenge shifts from technological capability to responsible orchestration. The systems that succeed will be those that amplify human purpose rather than replace it, turning autonomy into a collaborative fabric in which humans and intelligent agents co-create outcomes with measurable accountability and trust.

As autonomous systems mature, they are becoming one of the core building blocks of mass intelligence, translating human intent into coordinated, adaptive action across digital and physical environments.

Underlying concepts

Human-intent alignment

Autonomous systems act according to clearly defined human goals and ethics, amplifying human intent while keeping outcomes accountable and aligned with enterprise and societal priorities.

Transparent and explainable autonomy

Systems must explain their reasoning and actions in human-understandable terms, enabling trust, compliance, and effective collaboration between humans and intelligent agents.

Distributed human-machine collaboration

Humans and AI agents dynamically share tasks based on complementary strengths, forming cooperative networks where humans set direction and AI executes with adaptability.

Governance and ethical control frameworks

Autonomy requires continuous oversight, which encompasses regulatory compliance, ethical reviews, auditability, and the monitoring of system behavior and societal impact.

Adaptive autonomy across contexts

Autonomy levels adjust based on situational risk and required oversight, optimizing trust, safety and performance across business-critical and routine environments.

Integration of cognitive and physical autonomy

Digital intelligence and embodied systems — robots, vehicles, drones — interact with the physical world, connecting reasoning with real-world action.

Societal infrastructure of intelligence

Pervasive autonomy forms connected systems that support cities, industries and institutions, creating collective intelligence that serves human needs at scale.

Continuous human oversight and learning-feedback loops

Humans guide autonomous systems and learn from them in return, facilitating co-evolution, improved decisions and socially aligned outcomes.

Different forms of human oversight, from supervised decision support to self-regulating multiagent ecosystems, influence the way autonomy scales.

Autonomy orchestration stack



Figure 2: Autonomy orchestration stack

Technology landscape

The technology landscape of human-orchestrated autonomy is defined by a stack of capabilities that enable intelligent systems to act independently while remaining aligned with human oversight, organizational intent and accountable decision pathways.

● Now: Foundations of human-guided autonomy

Organizations are deploying the fundamental building blocks that allow autonomous systems to remain transparent, steerable and aligned with human purpose.

Agent frameworks and AI orchestration tools form the execution layer, enabling digital coworkers and task-specific agents to operate under human intent while improving decision speed. Surrounding this agentic core, explainability tools, model-monitoring frameworks and AI observability platforms generate the audit-grade visibility needed for responsible deployment. Human feedback pipelines and emerging AI governance platforms formalize supervision, treating agents as governed intelligent-automation assets, while hallucination management mechanisms help keep generative agents predictable, reducing automation errors.

Edge AI extends autonomy into physical environments — factories, vehicles, drones — bringing low-latency inference to where actions occur. Vision–language–action models and embodied AI capabilities create the bridge between digital reasoning and physical behavior, enabling human-supervised robots and autonomous devices to collaborate safely in real-world settings.

Sustainable AI practices ensure these systems scale within energy and computing constraints, aligning autonomy with sufficiency and long-term resilience.

● Next: Scaling coordination and autonomous decision layers

As agentic systems proliferate, enterprises begin to adopt the coordination capabilities required for autonomy at scale.

Multiagent orchestration frameworks emerge as the connective tissue that enables specialized agents to negotiate tasks, share context and escalate decisions. Autonomous decision-routing mechanisms dynamically determine when systems act independently and when human intervention is required, reinforcing safety and compliance. Transparent reasoning layers and neurosymbolic techniques add structure and verifiability to autonomous decision-making, helping organizations avoid “agent-washing” failures.

● Beyond: Toward system-level orchestrated intelligence

Over the longer horizon, autonomy becomes a system-level capability embedded in enterprises and society.

Autonomous governance protocols begin to operationalize policy-as-code, enabling self-regulating autonomous systems that can enforce constraints, detect violations and initiate corrective action. Collective-intelligence infrastructures coordinate distributed networks of humans and agents, forming organizational cognition that adapts continuously to context.

Research into bio-digital interfaces explores how human cognitive cues could eventually become part of the orchestration fabric, while adaptive learning societies extend autonomy into public services, mobility, healthcare and urban systems — transforming autonomy from a set of tools into a shared, intelligent societal infrastructure.

Technology radar

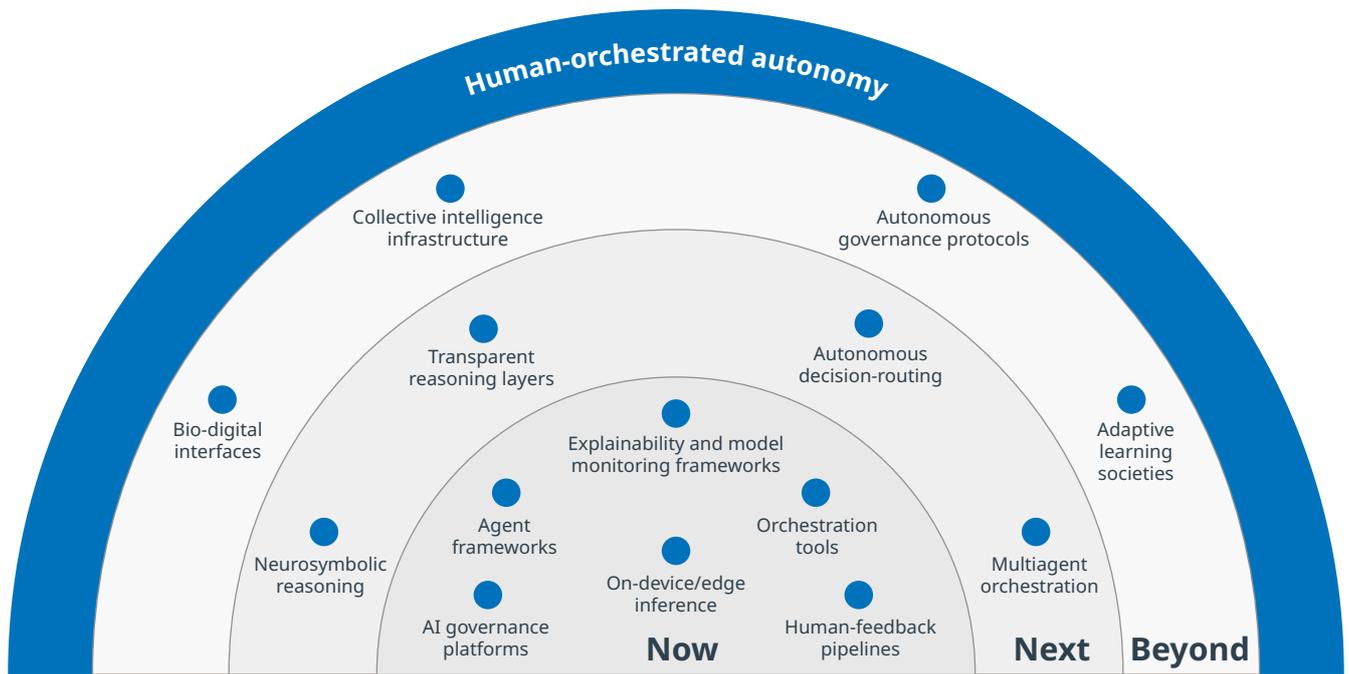


Figure 3: Human-orchestrated autonomy: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
Data orchestration	2.60	7.16	22.2%
Model orchestration	1.74	4.51	21.0%
GenAI orchestration	1.35	3.81	23.0%
Infrastructure orchestration	0.91	2.20	19.4%
Multiagent orchestration	0.63	1.85	24.1%

Table 1: Human-orchestrated autonomy: Market size and forecast ¹

Demand for orchestrated agents is expanding as enterprises shift from individual models to coordinated autonomy.

¹ Technavio, *AI Workflow Orchestration Market Analysis, Size, and Forecast 2025–2029*, July 2025.

Enterprise adoption signals

Gartner predicts up to

40%

of enterprise applications will include integrated task-specific agents by 2026, up from less than **5%** today. ²

By 2029, agentic AI will autonomously resolve

80%

of common customer service issues without human intervention, leading to a **30%** reduction in operational costs. ³

75%

of employees are comfortable working with AI agents, but only **30%** are comfortable being managed by them. ⁴

Industry partnership signals

The strongest momentum will be in the defense, mobility, manufacturing and logistics sectors, all of which require auditable autonomy under human oversight.

Adoption curve insights

2024 to 2026

Foundational agent frameworks, human-on-the-loop controls and autonomy guardrails become standard practice.

2026 to 2028

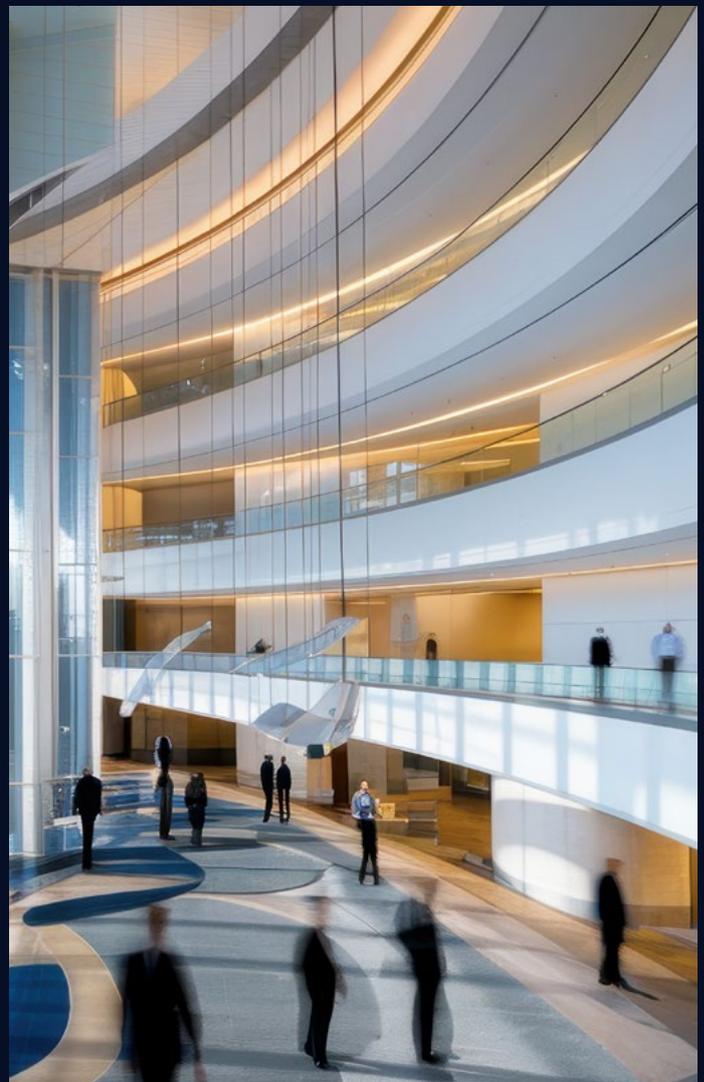
Multiagent orchestration, transparent reasoning layers and cognitive workflows scale across enterprise operations.

2028 to 2030

Autonomous decision-routing and coordinated human-machine ecosystems mature in critical functions.

2030+

System-level orchestrated intelligence integrates across sectors, forming adaptive, trust-aligned autonomous networks.



² Gartner. Press release. [Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025](#). August 26, 2025.

³ Gartner. Press release. [Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029](#). March 5, 2025.

⁴ Workday. Report. [AI Agents Are Here—But Don't Call Them Boss](#). © 2026.

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Arize AI	Machine learning (ML) observability and monitoring platform designed to detect, troubleshoot and explain model performance issues.	2020	Series C	131M
Character.AI	An AI companionship and conversational platform that lets users design and interact with personalized AI personas ("characters") that remember context and can be shared with others.	2021	Series A	193M
Credo AI	Provides an AI governance platform to ensure AI models are responsible, compliant and ethical.	2020	Series B	41.3M
Crew AI	Provides a Python-based open source multiagent framework and enterprise platform that lets teams build, orchestrate and deploy collaborating AI agents to automate complex workflows using various large language models (LLMs).	2024	Series A	18M
Fiddler AI	Offers an explainable AI (XAI) platform to monitor, explain and improve ML models for better transparency and fairness.	2018	Series B	63.8M
Labelbox	Platform for creating, managing and iterating on training data for ML applications.	2018	Series D	188.9M
LangChain	Open source framework designed to simplify the creation of applications using LLMs.	2022	Series B	160M
LlamaIndex	Data framework for connecting custom data sources to LLMs.	2022	Venture/ Series unknown	28.5M

Table 2: Human-orchestrated autonomy: Relevant startups ⁵

⁵ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate how human-orchestrated autonomy is already changing the way organizations operate, combining agentic systems with human oversight to deliver intelligence that is scalable, trustworthy and adaptive.

01 Autonomous decision support and supervisory workflows

Enterprises adopt human-supervised agents that draft documents and code, prepare decisions and resolve routine service requests, all while maintaining auditable, human-in-the-loop oversight.

Loan decisions, claims reviews, case-management summaries or eligibility determinations are generated autonomously, but every action is traceable: agent reasoning, source retrievals, confidence scores and escalation logic are all logged as part of a verifiable audit trail.

Early deployments show significant reductions in decision latency and error rates, particularly in supervised settings.

02 Human-supervised operational autonomy in physical and digital systems

Organizations deploy autonomous agents and robots that execute operational tasks — threat classification, lab analysis, inspection, predictive maintenance — under strategic human supervision.

In cyberdefense, agents detect and categorize threats before orchestrating automated containment steps; analysts intervene in high-risk scenarios. In scientific labs, autonomous analysis pipelines process data and propose actions for human verification. In manufacturing and logistics, embodied AI systems from companies like NEURA Robotics and Auterion coordinate multirobot operations with lower error rates when humans remain on the loop.

This is a shift from manual operation to orchestrated distributed intelligence, where intelligent agents execute the work and humans govern operational boundaries, safety rules and oversight thresholds.

03 Coordinated autonomy across physical ecosystems

Cities, grid operators and mobility providers adopt orchestrated autonomy to coordinate fleets, manage energy load and optimize real-time flows.

Autonomous vehicles and drones operate with human-on-the-loop supervision, while multiagent control systems manage routing, traffic conditions or mission priorities. Utilities apply the same principles to smart-grid balancing, where agents predict demand, reroute energy flows and prevent overload events, escalating exceptions to human operators. Urban mobility platforms orchestrate traffic, transit signals and autonomous assets, bringing together data, agents and human judgment into a unified decision network.

These ecosystems demonstrate how autonomy becomes a coordinating fabric, extending from digital workflows into the physical world — always with humans defining intent, constraints and ethical boundaries.

Future scenarios

AI control rooms become standard

Enterprises set up AI control rooms where teams monitor agent behaviors, hallucination risks and cognitive workflows, much like the way network operations centers (NOCs) and security operations center (SOCs) work today.



Primary uncertainties in this scenario

- Will organizations treat this as a compliance overhead or strategic advantage?
- Will real-time observability tools mature fast enough to support it?

Cognitive workflows replace process maps

Organizations shift from business process management (BPM)-style process automation to reasoning-chain automation, where AI agents negotiate tasks, escalate decisions and adapt workflows dynamically.



Primary uncertainties in this scenario

- Can cognitive frameworks (memory, planning) become reliable enough for core operations?
- Will businesses accept unpredictable but more adaptive workflows?

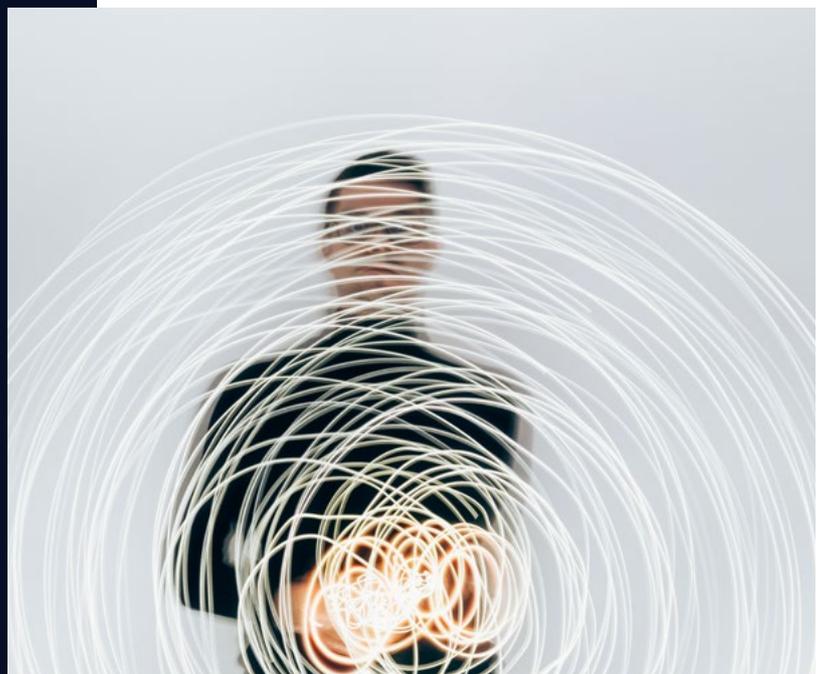
Autonomy constrained by human risk budgets

Enterprises assign each autonomous agent a “risk budget” that defines how much uncertainty it may act under — similar to financial risk limits.



Primary uncertainties in this scenario

- Will regulators mandate autonomy thresholds?
- Can we quantify and audit uncertainty well enough to operationalize this?



Strategic takeaways



Treat autonomy as a governed ecosystem, not a collection of agents

Organizations should design for orchestrated autonomy, ensuring that agents, models and workflows operate under clear supervisory rules, escalation paths and audit trails.

As enterprise applications with embedded agents become increasingly commonplace, leaders must establish governance frameworks early to avoid fragmented, high-risk deployments.



Invest in human-on-the-loop oversight as a core capability

The most successful deployments are those where humans remain strategic orchestrators. Supervised autonomy reduces error rates and prevents governance failures that contribute to the cancellation of agentic projects

Leaders should prioritize training, feedback loops, and new roles that manage autonomous systems rather than replace them.



Build autonomy into operational and physical systems to create measurable resilience

From cyberdefense to manufacturing and mobility, autonomy delivers significant performance gains when paired with low-latency edge inference and embodied AI.

Organizations should identify critical operations where autonomy can reduce decision latency and enhance continuity under stress.



- Are we designing autonomy to scale human intent — or to simply automate tasks?
- Are our people equipped to orchestrate, supervise and learn from autonomous systems rather than compete with them?
- Can we explain and justify autonomous decisions — to customers, regulators and society — when outcomes are contested or harmful?

2 Embodied agency and emotions

What it is and why it matters now

Systems that respond to data and the human experience

AI is no longer confined to text interfaces or isolated tasks but is moving into emotional, embodied and socially responsive forms. Digital humans, humanoid robots and affect-aware interfaces are becoming part of everyday life, from education and healthcare to mobility and public services. This shift emerges at a crucial time when societies face loneliness, skill shortages and rising demand for more inclusive, human-centric experiences.

Emotionally aware systems offer a pathway to trust, engagement and empathetic interaction, enabling machines to read gestures, tone and affect, and respond in ways that feel intuitive and supportive.

At a societal level, these capabilities transform service delivery and change expectations of digital systems, making synthetic empathy a structural driver of wellbeing, productivity and accessibility.

But as emotional AI becomes embedded in social infrastructure, its ethical design becomes urgent, requiring strong norms for privacy, authenticity and human agency.

By adding emotional awareness and embodied interaction to machine intelligence, this trend expands the relational dimension of mass intelligence, enabling systems that respond not only to data but also to the human experience itself.



Underlying concepts

Emotionally responsive systems as social infrastructure

Affective technologies evolve from isolated tools into foundational components of education, healthcare, mobility and governance, embedding empathy into digital ecosystems.

Sensorimotor empathy and affective feedback loops

Systems interpret multimodal cues — facial expression, posture, tone — and adjust responses in real time, building reciprocal feedback loops that enhance understanding and trust.

From emotional interaction to structural transformation

Affective intelligence changes the way services and relationships function, embedding emotion-driven design into societal, industrial and institutional systems.

Co-evolution of human and synthetic empathy

As artificial empathy matures, human emotional expectations shift, reshaping communication, learning and social connection.

Socioeconomic transformation through affective digitalization

Emotion-aware infrastructures support new models of inclusion, resilience and accessibility, making empathy a driver of social and economic progress.

Embodied emotional presence in human environments

Synthetic agents use voice, movement and gestures to convey presence and empathy, creating intuitive shared spaces across physical and virtual environments.

Synthetic emotion as a trust and communication catalyst

Simulated emotion improves clarity and intent signaling, strengthening long-term trust in personal, organizational and civic interactions.

Emotionally intelligent work and collaboration ecosystems

Emotional awareness enhances productivity and wellbeing in hybrid work settings, supporting an empathetic work culture and more sustainable digital workplaces.

Ethical and societal design of emotional systems

Emotional AI demands strong norms for privacy, transparency and agency to ensure inclusivity, safety and integrity in emotionally responsive systems.

The move to embodied agents brings emotion to the foreground, transforming interaction from a functional exchange to relational engagement.

Embodied interaction spectrum

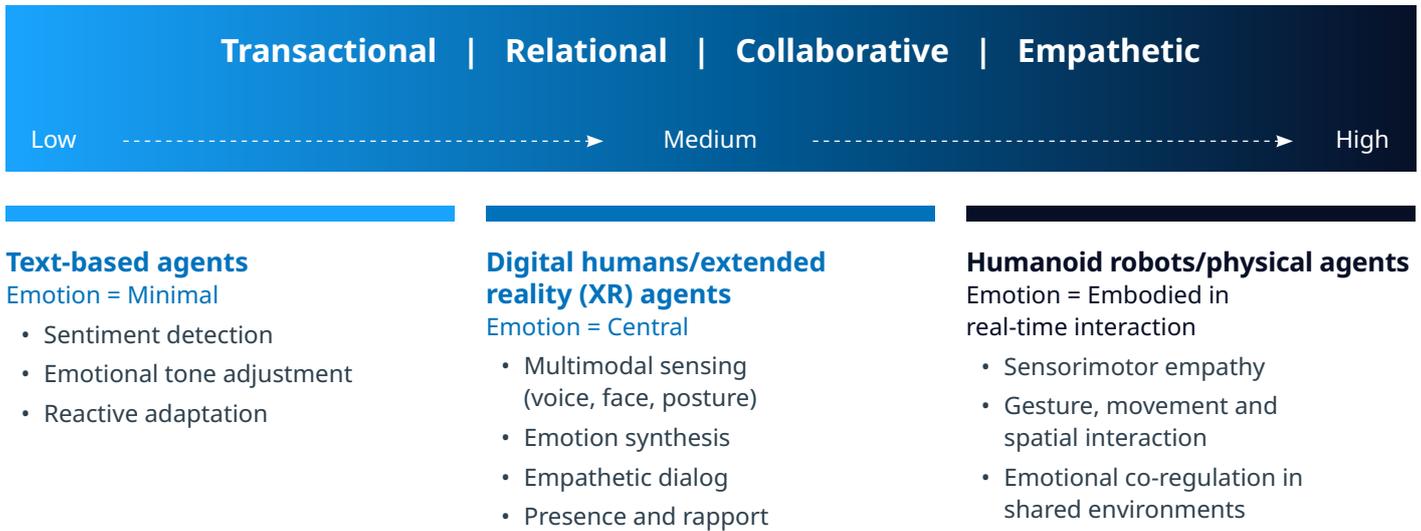


Figure 4: Embodied interaction spectrum



Technology landscape

Embodied agency and emotions is enabled by a rapidly maturing stack of multimodal sensing, expressive interfaces and synthetic emotional intelligence. These technologies shift AI from cognitive interaction to relational, empathetic systems that can perceive, interpret and respond to human affect across contexts.

● **Now:** Mature capabilities forming today's emotional infrastructure

Current deployments rely on affective computing frameworks, emotion recognition, real-time sentiment analysis and multimodal affect sensing, which improve accuracy over unimodal models.

Clinically validated wearable affect sensors (electrodermal activity (EDA) and heart rate variability (HRV)) deliver up to 84% accuracy in detecting stress, enabling healthcare, wellbeing and mobility applications.⁶

Emotionally responsive interaction is supported by speech–emotion synthesis, advanced virtual assistants, AI companions, and adaptive learning and wellbeing platforms that reduce support time and improve retention and wellbeing outcomes.

Physical embodiment is provided through social and care robotics and haptic feedback devices, while deepfake protection and affective safety mechanisms ensure emotional integrity and safeguard against manipulation.

● **Next:** Scaling technologies enabling embodied emotional interaction

The next wave brings deeper embodiment and richer multimodal intelligence.

Embodied AI, embodied conversational agents and sensorimotor empathy frameworks enable systems that respond not only to emotional signals but also to movement, posture and context.

Digital humans evolve into emotionally expressive, adaptive interfaces. Synthetic emotion engines, affective-twin systems and empathetic copilot frameworks introduce longitudinal emotional modeling and context-aware adaptation.

At the intelligence layer, vision–language–action models unify perception, communication and action, enabling the emotionally informed orchestration of digital and physical tasks.

● **Beyond:** Frontier technologies creating autonomous affective environments

Longer-term development points to systems capable of autonomous emotional adaptation.

Embodied generalist agents combine multimodal perception, action and affective reasoning. Neuroadaptive interfaces adjust environments or interactions based on cognitive and emotional load.

Synthetic empathy coprocessors bring low-latency emotional inference to edge devices, and autonomous affective environments — from workplaces to healthcare settings — continuously adapt lighting, sound and interaction flows based on collective emotional states.

⁶ Kyriakou K, Resch B, Sagl G, et al. *Detecting Moments of Stress from Measurements of Wearable Physiological Sensors*. Sensors (Basel). 2019;19(17):3805. Published 2019 Sep 3. doi:10.3390/s19173805

Technology radar

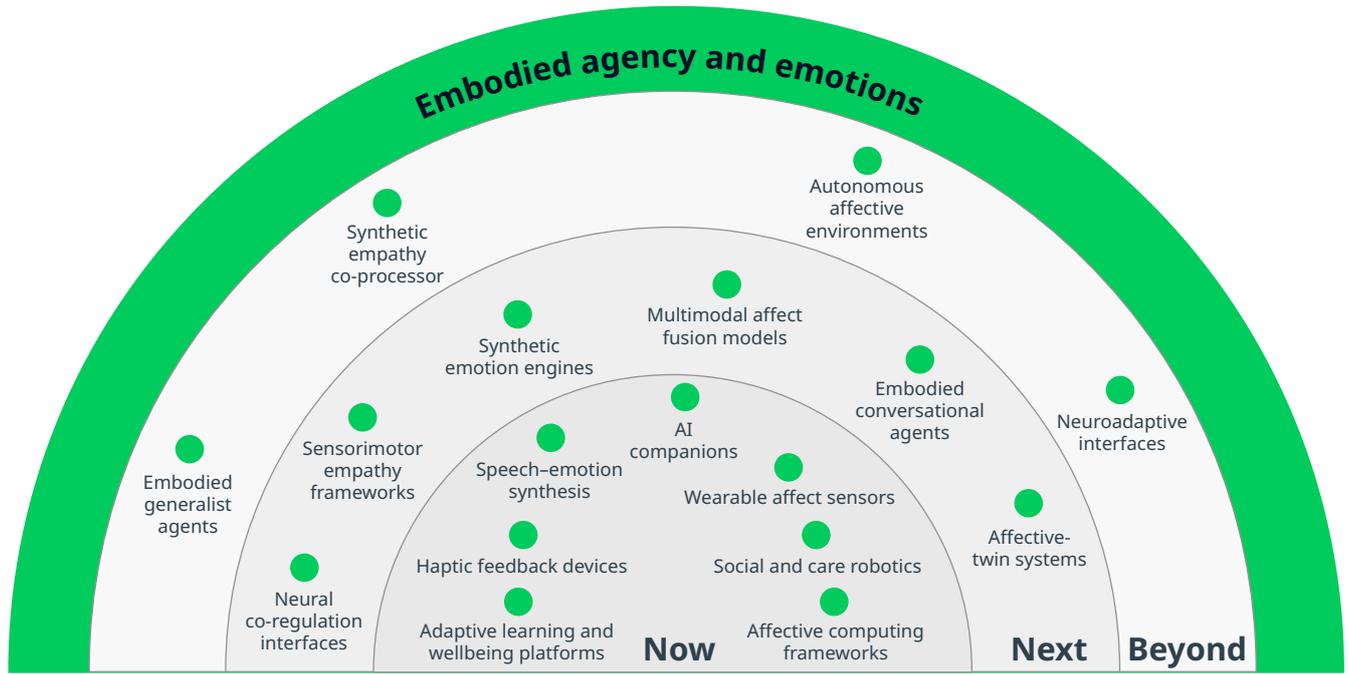


Figure 5: Embodied agency and emotions: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
Global affective computing	88.57	347.43	31.4%
Global emotion AI	5.99	17.42	23.8%
Global digital human avatar	5.29	32.72	43.9%
Global emotion analytics	3.91	9.26 *	18.2%

Table 3: Embodied agency and emotions: Market size and forecast ⁷

*Internal estimates

Emotion AI and digital humans grow fast as industries embed affect-aware interaction in service models.

⁷ Sources: Technavio reports. Global Affective Computing Market 2025–2029; Global Emotion AI Market 2025–2029; Global Digital Human Avatar Market 2025–2029; Global Emotion Analytics Market 2024–2028

Enterprise adoption signals

In 2025, emotion AI is described as “core,” with

88%

of marketers using it in their daily work.⁸

Customer experience monitoring is already the top application area for emotion AI, accounting for over

31%

of total demand.⁹

Patient engagement accounts for

62%

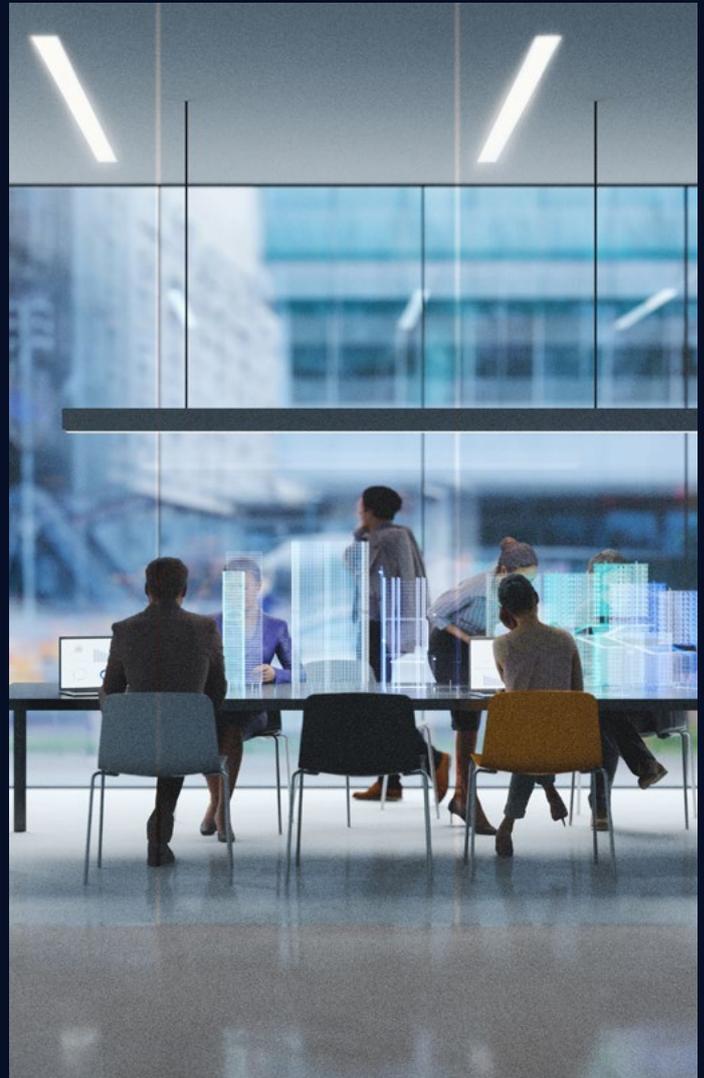
of healthcare AI use (including digital humans), while administrative efficiency accounts for **60%**.¹⁰

Industry partnership signals

The highest activity is in healthcare, automotive safety, education, customer experience (CX) and elder care — sectors that depend on emotional accuracy and trust.

Adoption curve insights

- **2024 to 2026**
Multimodal emotional sensing becomes the technical standard.
- **2026 to 2028**
Digital humans, extended reality (XR) agents and emotion-adaptive interfaces scale into production.
- **2028 to 2030**
Emotion AI becomes part of our social infrastructure, embedded in care, mobility, education and CX.
- **2030+**
Synthetic empathy and human-AI co-regulation reshape societal interaction models.



^{8,9} Electro IQ. *Emotion AI Statistics by Market Size, Segmental Analysis, Impact, Usage and Facts (2025)*, November 29, 2025.

¹⁰ Market.US. *Global AI-Powered Digital Humans Market Size, Share and Analysis Report*, December 2025.

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Empatica	Develops digital biomarkers and AI systems using wearable sensors for health monitoring, especially for epilepsy and stress.	2011	Series B	35.6M
Furhat Robotics	Designs and develops social robots with human-like features and conversational AI capabilities.	2014	Seed	9.8M
HeyGen	Provides AI tools to generate studio-quality avatar videos, multilingual lip-synced translations, and personalized clips for marketing, sales, training and internal communications.	2020	Series A	74M
Hume AI	Focused on building empathetic AI using models trained on human emotional and vocal expressions.	2021	Series B	72.8M
Intuition Robotics	Develops empathetic digital companions and social robots like ElliQ to improve the lives of older adults.	2016	Venture/Series unknown	83M
LuxAI	Specializes in disruptive robotic solutions for education, healthcare and entertainment.	2018	Grant	71.4K
Ultraleap	Developer of hand-tracking and haptic-feedback technologies for immersive computing (augmented reality/virtual reality (AR/VR); metaverse).	2013	Series D	166.8M

Table 4: Embodied agency and emotions: Relevant startups ¹¹

¹¹ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate the emergence of embodied agency and emotions in sectors where human wellbeing, trust and high-quality interaction are critical.

01 Emotion-adaptive care and wellbeing assistants

Emotionally responsive agents support clinicians and caregivers by continuously sensing stress, anxiety, or cognitive load through multimodal signals (voice, facial cues, EDA/HRV).

Such agents include care robots, therapeutic assistants and wearable affect-tracking systems. Clinical evidence shows affective wearables reduce monitoring workloads, while emotionally aware virtual agents decrease patient support time and increase treatment adherence.

These systems augment workforce capacity, enable early intervention and provide empathetic companionship in elder care, chronic-condition management and mental-health support environments.

02 Emotionally aware enterprise interaction and collaboration

Next-generation enterprise tools integrate emotion recognition into everyday collaboration.

Emotionally aware meeting facilitators detect disengagement, frustration or burnout risk, adjusting interaction flow or offering supportive prompts. Affective leadership coaches provide managers with real-time emotional insights to improve communication, reduce conflict and strengthen team wellbeing, thereby decreasing burnout risk in hybrid teams.

In education and training, emotionally adaptive tutors respond to confusion or stress, improving engagement and learning outcomes.

03 Emotion-driven service and experience personalization

Emotionally aware systems transform service journeys by responding to user affect in real time.

Affective service recovery identifies frustration early, enabling digital humans or AI agents to adjust tone, escalate proactively or tailor solutions.

Driver-emotion monitoring is becoming a key safety feature in mobility, with most new vehicles expected to include such sensing capabilities.

Meanwhile, entertainment and consumer engagement evolve through personalized, emotionally responsive entertainers and companions that use synthetic emotion engines and adaptive expression to create immersive, empathetic or supportive experiences.

Future scenarios

Affective metrics enter workplace dashboards

AI systems begin tracking team sentiment, stress and emotional climate in real time, feeding these metrics into HR or operational dashboards.



Primary uncertainties in this scenario

- Will employees accept being emotionally monitored?
- How will privacy regulators treat affective signals?

Digital humans become operational gatekeepers

Customer-facing digital humans, chosen for emotional stability and consistency, take over onboarding, service triage, coaching and conflict de-escalation.



Primary uncertainties in this scenario

- How will customers respond to emotionally competent AI personas?
- Can emotional models maintain consistency across modalities and contexts?

Behavioral-influence systems face pharmaceutical-level regulation

Governments classify certain types of affective nudging or influence engineering as regulated "behavioral interventions."



Primary uncertainties in this scenario

- Who decides what constitutes harmful influence and beneficial influence?
- Can global alignment on affective AI regulation be achieved?



Strategic takeaways



Build emotional intelligence into every human-facing system

Emotion recognition, multimodal sensing and emotionally adaptive interfaces are no longer experimental. They are becoming expected features in healthcare, education, customer service and enterprise collaboration.

Organizations should start by integrating affective signals into existing products and workflows to create measurable improvements in engagement, satisfaction and performance, starting with low-risk use cases such as service recovery.



Treat emotion data as a high-sensitivity asset

Emotion signals — voice tension, facial cues, biometrics, microexpressions — carry psychological and behavioral implications that require stricter governance than what's required for traditional personal data.

Companies should establish emotional-data policies early. These include clear consent protocols, transparency requirements, bias-mitigation processes and safeguards against manipulation. Emotional integrity will become a key differentiator in trust-centric sectors.



Prepare for emotionally intelligent agents that co-work, co-teach and co-care

Digital humans, affective copilots and embodied assistants will soon operate alongside employees in coaching, facilitation, wellbeing and frontline- service roles.

Organizations should begin designing hybrid workflows, training managers to interpret emotional insights, and defining role boundaries between human and synthetic empathy.

Early adopters will gain significant advantages in workforce resilience, customer engagement and service personalization.



- How will our organization ensure that emotionally intelligent systems enhance human wellbeing and agency rather than influence or automate emotions in ways we cannot fully oversee or explain?
- Where do we draw the line between empathetic support and emotional manipulation, and who is accountable for enforcing it?
- Are we prepared to treat emotion and affective data as high-sensitivity assets across their entire lifecycle?

3 Intelligence we trust



What it is and why it matters now

The stabilizing pillar of mass intelligence

As AI systems become more autonomous, interconnected, and embedded in critical operations, trust becomes the essential operating layer of digital civilization. Organizations face increasingly adaptive cyberthreats, rapid growth in AI-generated misinformation and expanding reliance on opaque models. At the same time, autonomy and emotional capabilities emerging across industries amplify the stakes of integrity, security and transparency.

Beyond enterprise governance, trustworthy intelligence becomes a societal foundation, supporting transparency, accountability and confidence at scale.

This requires a shift from conventional cybersecurity to adaptive, self-learning, proactive defense ecosystems that safeguard both data and machine intelligence itself.

Zero trust architectures, cognitive transparency and verifiable data provenance are no longer best-practice ideals — they are prerequisites for safety, regulation and societal confidence.

The systems that thrive will be those that make intelligence explainable, protected, and aligned with human values, ensuring that progress is built on a foundation of reliability and ethical assurance.

Underlying concepts

Adaptive and resilient AI systems

Self-assessing, anomaly-detecting systems maintain performance under uncertainty, ensuring continuity in operational landscapes that face dynamic threat patterns.

Cognitive transparency and explainable intelligence

Understanding how AI reaches conclusions enables verification, accountability and predictable system behavior.

Zero trust architectures

Every user, device and algorithm is continuously authenticated and validated, minimizing implicit trust and reducing attack surfaces.

Authenticity and alignment of machine knowledge

Data, models and inferences must remain tamper-proof and aligned with human objectives, supported by mechanisms that detect drift, poisoning or misinformation.

Active cyberdefense and predictive security

AI anticipates attacks, simulates adversarial behavior and orchestrates adaptive countermeasures.

Data integrity and provenance assurance

Cryptographic verification and lineage tracking preserve the origin and consistency of information across systems.

Governance and ethical oversight of intelligent systems

Policies and institutional structures embed fairness, accountability and compliance into AI operations.

Human-machine collaboration and trust calibration

Transparent interfaces help humans calibrate reliance on AI, balancing confidence with oversight.

Social and institutional resilience through trust networks

Shared intelligence networks, standards and cross-organizational frameworks strengthen collective assurance and resilience.

Cultural and psychological dimensions of digital trust

Understanding societal expectations supports the design of systems perceived as fair, reliable and controllable.

Trustworthy intelligence is a stabilizing pillar of mass intelligence, ensuring that increasingly autonomous, connected and adaptive systems remain transparent, secure and aligned with human values.

Trust flow pipeline



Figure 6: Trust flow pipeline



Technology landscape

The technology landscape of intelligence we trust comprises a stack of capabilities that make intelligent systems transparent, verifiable, and secure by design.

● Now: Core trust foundations

Organizations are deploying foundational controls that ensure integrity and accountability across AI systems.

These include explainable AI (XAI), model interpretability tools, AI audit and compliance platforms, and data lineage and provenance systems. Together, these enable audit-grade evidence for regulatory readiness. Security baselines such as zero trust security, privacy-centric identity infrastructures and homomorphic encryption for AI protect sensitive pipelines. With the rise of multimodal manipulation, Coalition for Content Provenance and Authenticity (C2PA) content authenticity, model chain-of-trust systems and deepfake protection have become essential safeguards. Adversarial robustness is strengthened through AI red teaming, while hallucination management keeps generative systems predictable and aligned.

● Next: Scaling governance and runtime assurance

As AI becomes more autonomous, enterprises expand into capabilities that operationalize trust at scale.

AI observability and auditability platforms provide continuous monitoring of drift, lineage and inference behavior, while agent governance and compliance frameworks define safe operating boundaries for AI agents. Governance workflows evolve through AI governance, risk and compliance (GRC) automation, policy-compliance agents, and self-validating models that can detect and correct inconsistencies. To secure long-term resilience, organizations integrate quantum-safe encryption, supported by emerging cross-organizational trust networks that share signals and threat intelligence.

● Beyond: System-level trust infrastructure

Looking ahead, trust extends beyond individual enterprises to shared digital infrastructures.

Global AI trust registries will verify model identity, provenance and compliance across organizational borders, while cognitive traceability protocols standardize how decisions are recorded in complex AI ecosystems. In the societal context, AI misinformation firewalls will act as protective layers against sophisticated synthetic-media and influence attacks.

Technology radar

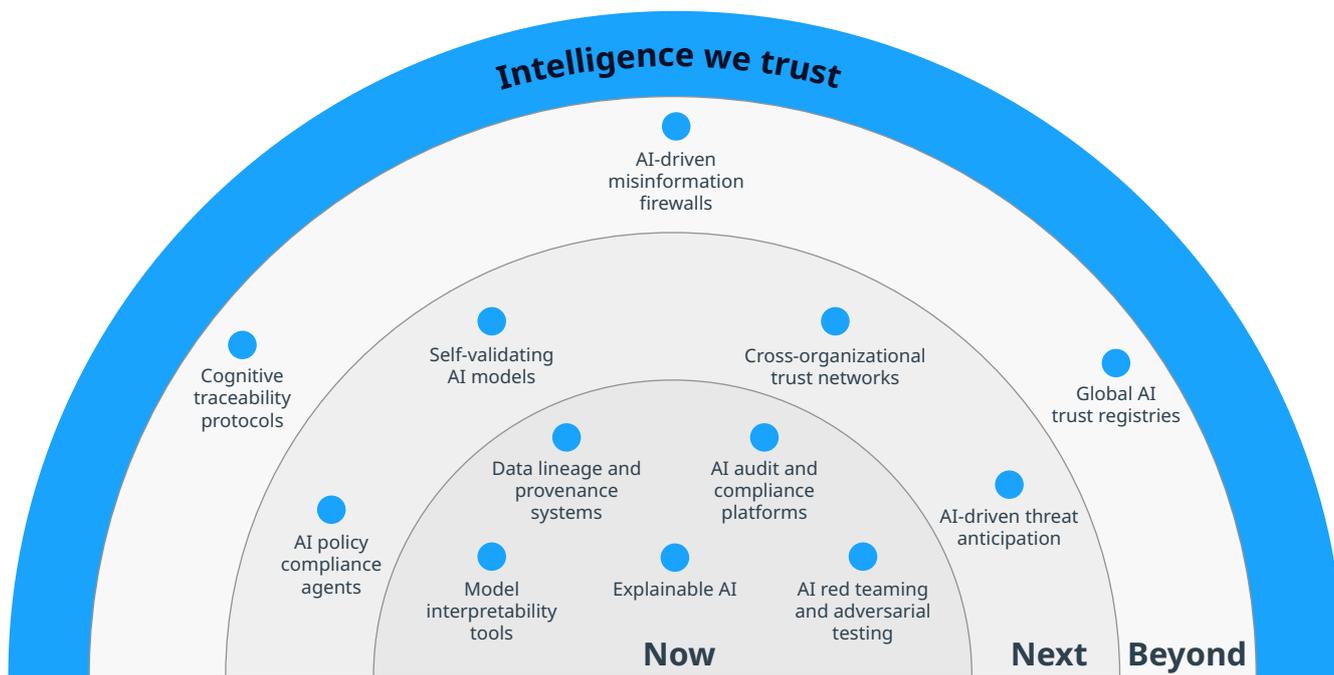


Figure 7: Intelligence we trust: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
AI trust, risk and security management (TRISM)	2.34	6.22 *	21.6%
AI observability	1.20	3.60 *	24.6%
AI in cybersecurity	28.00 *	79.80 *	21.7%

Table 5: Intelligence we trust: Market size and forecast ¹²

*Internal estimates

Trust functions are converging into unified stacks spanning governance, observability, provenance and cybersecurity.

¹² Sources:

GlobalNewswire. 2025. [AI Trust, Risk and Security Management Trends Analysis and Growth Forecasts, 2025-2030: Market to Reach \\$7.44 Billion - Generative AI Foundation Models Spur Demand for Advanced AI TRISM Tools](#). 21 May, 2025.

Market Intelo. [AI Observability Market Research Report 2033](#), 2025.

Fortune Business Insights. 2026. [Artificial Intelligence in Cybersecurity Market Size, Share & Industry Analysis](#). February 9, 2026.

Enterprise adoption signals

By 2028, precipitated by AI bill of materials (BOM) requirements,

85%

of data products will include a data BOM detailing data collection, edits made, data cleanup, and how consent was obtained.¹³

90%

of enterprises will be adopting agent governance frameworks within **18 months**.¹⁴



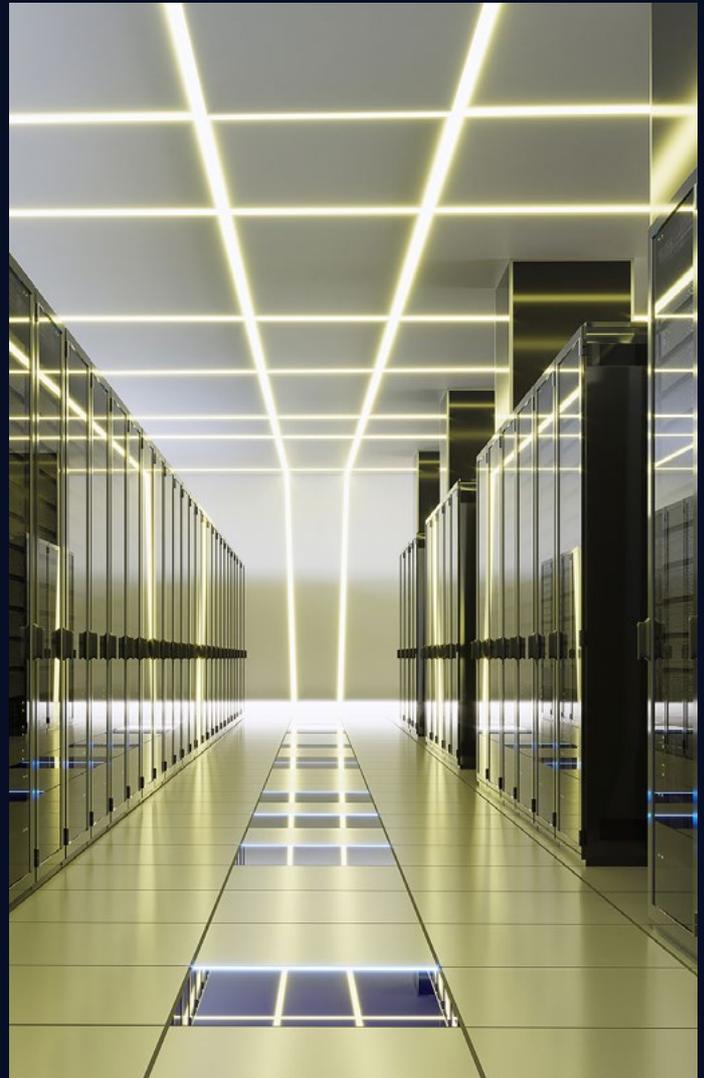
Continuous monitoring replaces periodic validation; **trust stacks emerge as standard enterprise architecture**.

Industry partnership signals

Intelligence we trust is strongest in financial services, healthcare, public sector and digital platforms, all of which require verifiable AI behavior and provenance.

Adoption curve insights

- **2024 to 2026**
Evidence-producing governance, provenance and continuous monitoring become baseline requirements.
- **2026 to 2028**
Trust stacks — combining security, observability and lineage — are widely institutionalized across enterprises.
- **2028 to 2030**
Trust becomes operational infrastructure, enabling verifiable AI decisions in regulated and safety-critical domains.
- **2030+**
Federated trust networks emerge, supporting cross-organizational validation and global AI.



¹³ IDC. 2024. [IDC FutureScape: Worldwide Security and Trust 2025 Predictions](#). October 2024.

¹⁴ Everest Group. Webinar. [AI Safety: The Next Frontier in Trust & Safety](#). October 28, 2025.

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Alma Security	Focuses on cloud-native application protection platforms (CNAPPs) for securing cloud environments and applications.	2022	Seed/Early stage	5.5M
Aporia	Platform for monitoring, explaining and validating machine learning models in production.	2019	Series A (Acquired)	30M
FairNow	Focuses on automated compliance and risk management for the responsible use of artificial intelligence and machine learning (AI/ML) in enterprises.	2023	Seed/Early stage	3.5M
Qodo	Provides tools and platforms for data governance and compliance.	2022	Series A	50.6M
Scribe Security	Provides a security and compliance platform for the software supply chain, focusing on software bills of materials (SBOMs).	2021	Seed/Early stage	7M
Theom	A data-access governance platform that uses graph-based security models to protect data across cloud environments.	2020	Series A	36.4M
WitnessAI	Provides tools for auditing, evaluating and ensuring the safety and ethical use of LLMs.	2023	Non-equity assistance	27.5M
Xapien	Offers an AI-powered research platform for due diligence, compliance and background checks.	2018	Series A	18.5M

Table 6: Intelligence we trust: Relevant startups ¹⁵

¹⁵ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate how the rise of intelligence we trust enables new cross-industry applications that combine governance, provenance, explainability and adaptive security.

01 Continuous AI assurance in regulated workflows

Enterprises directly integrate AI audit and certification platforms, model provenance, version tracking and explainability layers into decision systems.

Every model decision — loan approval, medical triage recommendation or eligibility assessment — produces audit-grade evidence, including lineage, consent and model-version metadata.

Predictive flags signal drift or fairness deviations, while red team evaluations run automatically in the background. AI products are trending to requiring audit-grade training-data evidence, underscoring the fact that continuous assurance is becoming a structural requirement rather than a compliance add-on.

02 Adaptive threat anticipation and secure AI operations

Organizations deploy predictive threat intelligence, behavioral trust analytics, predictive threat analytics and automated vulnerability detection to secure AI ecosystems.

Models continuously learn from adversarial patterns, simulate attacks, and detect anomalies in user, agent and system behavior.

This is a shift from reactive defense to active cyberdefense and predictive security, where AI anticipates threats, orchestrates countermeasures and validates data integrity through provenance functions.

03 Trusted knowledge and fact-verification assistants

Enterprises deploy fact-verification assistants integrated with explainability layers, content provenance and C2PA authenticity signals to secure information flows.

These assistants validate claims, detect synthetic or manipulated content, and generate traceable citations for internal and external communications.

Provenance verification and synthetic-media detection are critical to maintain authenticity, transparency and societal trust across interconnected ecosystems.

Future scenarios

Provenance-verified content becomes a baseline feature

Most enterprise content (reports, decisions, data pipelines) requires cryptographic provenance stamps (C2PA, chain-of-trust).



Primary uncertainties in this scenario

- Will adoption be driven by regulation or competitive pressure?
- Can provenance verification scale without slowing operations?

Trust scores for AI govern procurement decisions

Enterprises purchase AI systems based on standardized trust metrics: robustness, transparency, the risk of misuse risk and the integrity of provenance.



Primary uncertainties in this scenario

- Who will define global trust standards? ISO? Governments? Industry alliances?
- Can trust scores meaningfully reflect real risk or will they be gamed?

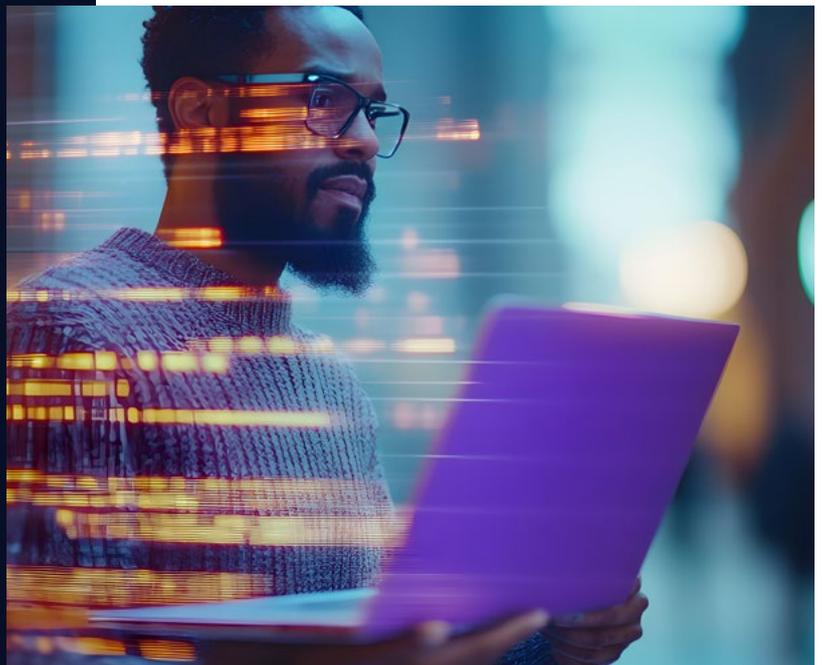
Quantum-safe identity and data integrity break vendor lock-in

Quantum-safe identity systems enable AI agents, data, and governance profiles to move securely between clouds and sovereign ecosystems.



Primary uncertainties in this scenario

- Will global quantum-safe standards converge or fragment geopolitically?
- Can vendors resist market pressure to keep identity proprietary?



Strategic takeaways



Make trust a runtime capability, not a compliance layer

Organizations must shift from periodic audits to continuous AI assurance, integrating explainability, provenance, drift monitoring and automated evidence-generation directly into AI pipelines. Trust becomes an operational requirement that determines whether AI can scale safely across high-stakes environments.



Strengthen AI security before scaling autonomy

As AI systems become more agentic, enterprises must strengthen them with predictive threat intelligence, behavioral trust analytics and automated vulnerability detection. Protecting models, data, and agentic workflows from poisoning, manipulation and synthetic-media risks is essential to preserve system integrity.



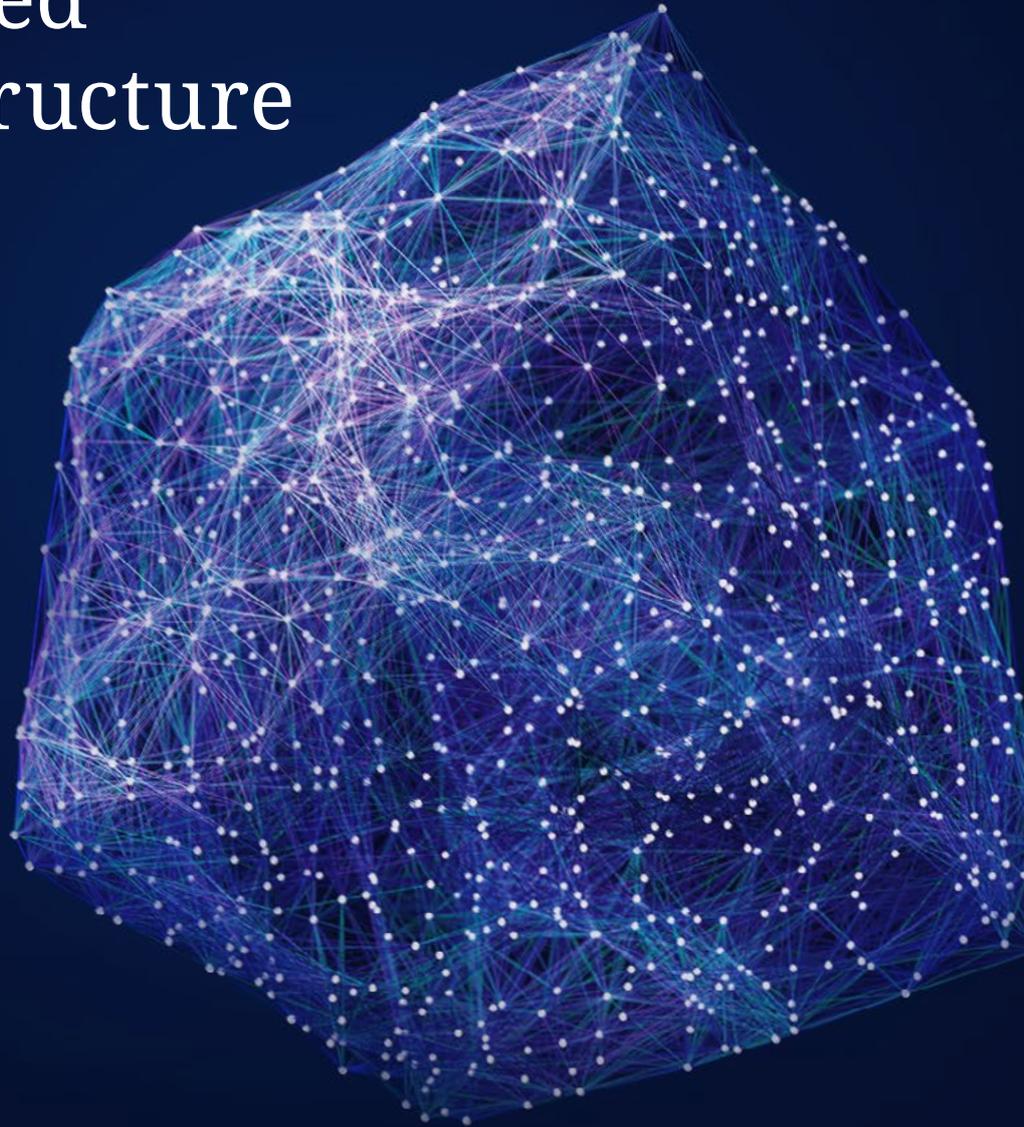
Build cross-organizational trust networks

No organization can defend itself against evolving AI risks alone. Shared provenance standards, content-authenticity frameworks and trust-intelligence exchanges are emerging as critical infrastructure, enabling industries to coordinate defenses and accelerate their regulatory readiness.



- If every AI decision in our organization had to be independently verified tomorrow, how much of our current AI estate would be able to produce trustworthy, auditable evidence on demand?
- Can our trust mechanisms keep pace with adaptive threats, model drift and AI-generated manipulation?
- Are our governance frameworks designed for individual models or for ecosystems of agents, data and decisions acting together?

4 Informed infrastructure



What it is and why it matters now

The operational substrate of mass intelligence

Digital infrastructure is undergoing a profound transformation. Infrastructure layers that were once passive back-end systems are becoming cognitive, anticipatory and sovereign, changing the way organizations manage workloads, resources and resilience. This comes at a time when demand for computation is exploding — driven by AI, connected devices and distributed intelligence — while energy, sovereignty and sustainability constraints intensify.

Informed infrastructure provides a strategic advantage: the ability to sense, model, simulate and optimize entire environments before decisions are made.

High-performance computing, edge-to-cloud continuity and quantum-enabled simulation enable governments and enterprises to transition from reactive operations to predictive, resilient and transparent digital foundations. As nations pursue autonomy over data and computation, informed infrastructure becomes the backbone of both innovation and sovereignty.

It also provides the operational substrate of mass intelligence — an environment where computing, connectivity and sensing continuously learn and optimize to support intelligent activity at every scale.

Underlying concepts

Cognitive infrastructure systems

AI, IoT and continuous intelligence turn infrastructure into adaptive systems that sense, learn and optimize in real time.

AI-augmented decision ecosystems

Hybrid analytics and reasoning engines support real-time, evidence-based decisions, balancing agility, cost and sustainability.

Sovereign digital foundations

Transparent, secure and interoperable ecosystems give nations and enterprises control over data, computation and infrastructure governance.

Interconnected digital-physical continuum

Digital twins and sensors link physical systems with virtual models, enabling advanced simulation, predictive operations and safer environments.

Resilient and anticipatory infrastructure design

Predictive modeling identifies stress points early, embedding resilience into planning and operations.

High-performance computing and quantum-enhanced simulation

Advanced computation models complex systems — cities, grids, logistics, materials — enabling virtual testing before making real-world decisions.

Seamless hybrid infrastructure

Cloud, edge and devices interoperate as a unified environment, enabling optimization from sensor-level to strategic control.

Sustainable systemic optimization

Continuous intelligence enables the optimization of emissions, resources, equity and efficiency as measurable operational outcomes.

Adaptive governance and trust frameworks

Transparent oversight, participation and ethical AI guardrails build trust in automated infrastructure.

Sociotechnical integration for human-centric outcomes

Infrastructure succeeds when it enables equitable access, inclusion and human wellbeing.

At the heart of informed infrastructure is a distributed model where cloud-scale intelligence, local inference and on-device sensing operate as a coordinated whole.

Layers of informed infrastructure

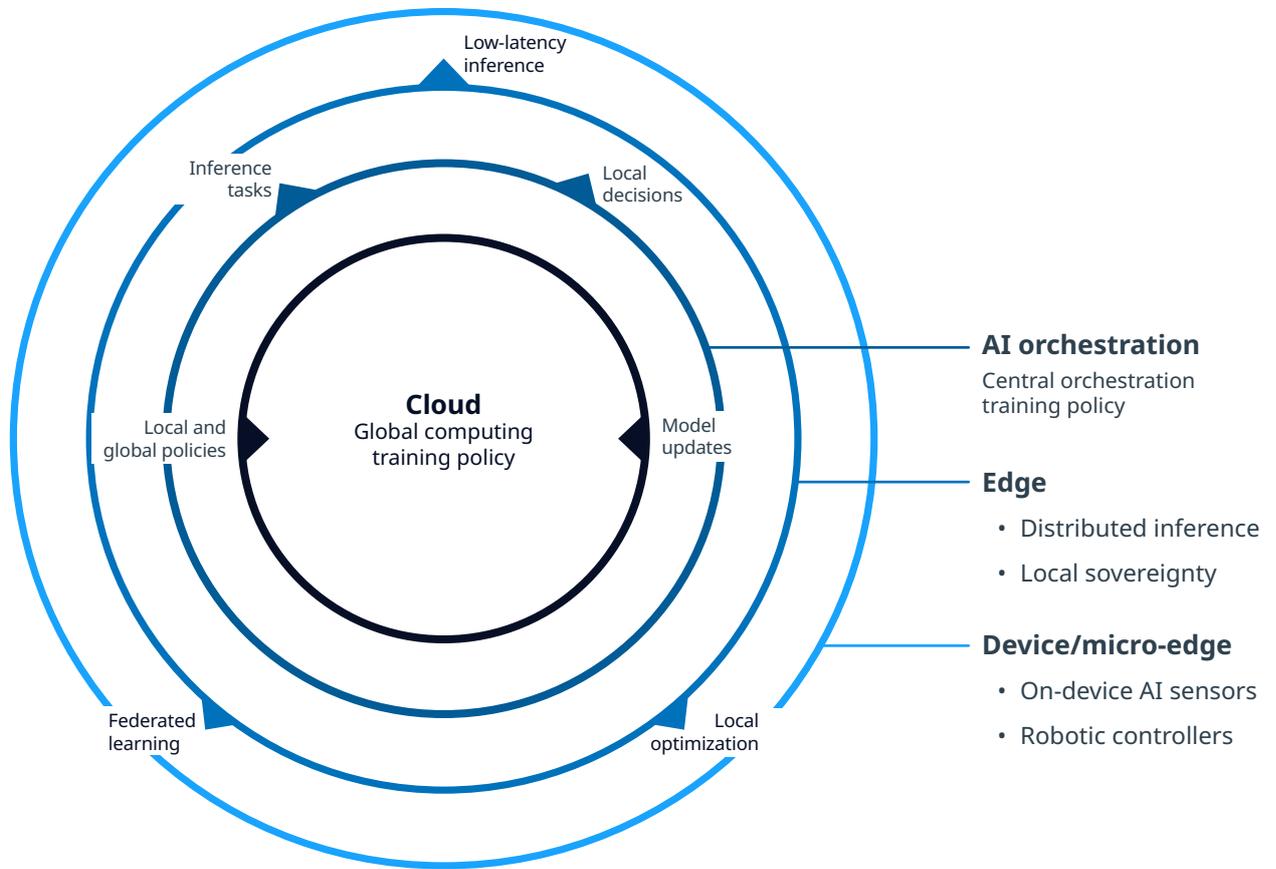


Figure 8: Layers of informed infrastructure



Technology landscape

Informed infrastructure emerges at the intersection of AI-optimized computing, hybrid cloud-edge architectures and sovereign digital foundations. Its technology landscape evolves across three horizons.

● **Now: Foundations of intelligent infrastructure**

Organizations are strengthening the core layers of their digital ecosystems as they prepare for an AI-intensive operating environment.

This shift is accelerating as a growing share of workloads demand real-time, edge-level intelligence, while only a minority of enterprises have modernized their infrastructure enough to support AI at scale.

In this context, cloud-edge orchestration tools become essential for distributing workloads intelligently, and local sovereign-cloud initiatives ensure jurisdictional control and compliant data-handling. Open data and interoperability standards further enable seamless integration across hybrid systems. Together, these technologies establish the foundational environment for connected, compliant and responsive infrastructure that can support emerging AI-driven operations.

● **Next: Scaling transparency, sustainability and distributed intelligence**

As AI penetrates deeper into operations, infrastructure must evolve to support greater transparency, resilience and distributed intelligence.

Self-healing and anticipatory infrastructure introduces predictive resilience, improving reliability and reducing unplanned downtime. This is complemented by hybrid cloud-edge-device orchestration, which enables workloads to run optimally across distributed physical and digital environments.

To protect data while expanding intelligence at the edge, federated learning and privacy-preserving analytics allow organizations to derive insight where data is created, without compromising compliance. Meanwhile, sustainable computing infrastructure accelerates as enterprises adopt energy-aware operations, demonstrating that sustainability is increasingly embedded in infrastructure design rather than treated as an external reporting requirement.

Overseeing this complexity, adaptive governance engines introduce explainability, auditability and policy automation across AI-driven systems. Collectively, these capabilities elevate infrastructure from reactive management to proactively coordinated, distributed intelligence with measurable operational improvements.

● **Beyond: Toward autonomous, sovereign infrastructure ecosystems**

Looking further ahead, infrastructure evolves into a highly autonomous, sovereign-aligned and globally interoperable fabric.

Cognitive infrastructure layers integrate AI, high-performance computing and sensor networks to support continuous prediction and optimization across systems. As these capabilities mature, fully autonomous infrastructure ecosystems emerge, coordinating computing, energy and data flows with minimal human intervention.

Beyond individual organizations, interoperable global infrastructure frameworks will enable cross-border collaboration on computing resources, sustainability practices and digital governance. At the same time, sovereign computing networks provide nations and regions with durable control over critical AI capabilities, ensuring long-term digital resilience. These trajectories point to infrastructure becoming an intelligent, strategic operating layer for economies and societies.

Technology radar

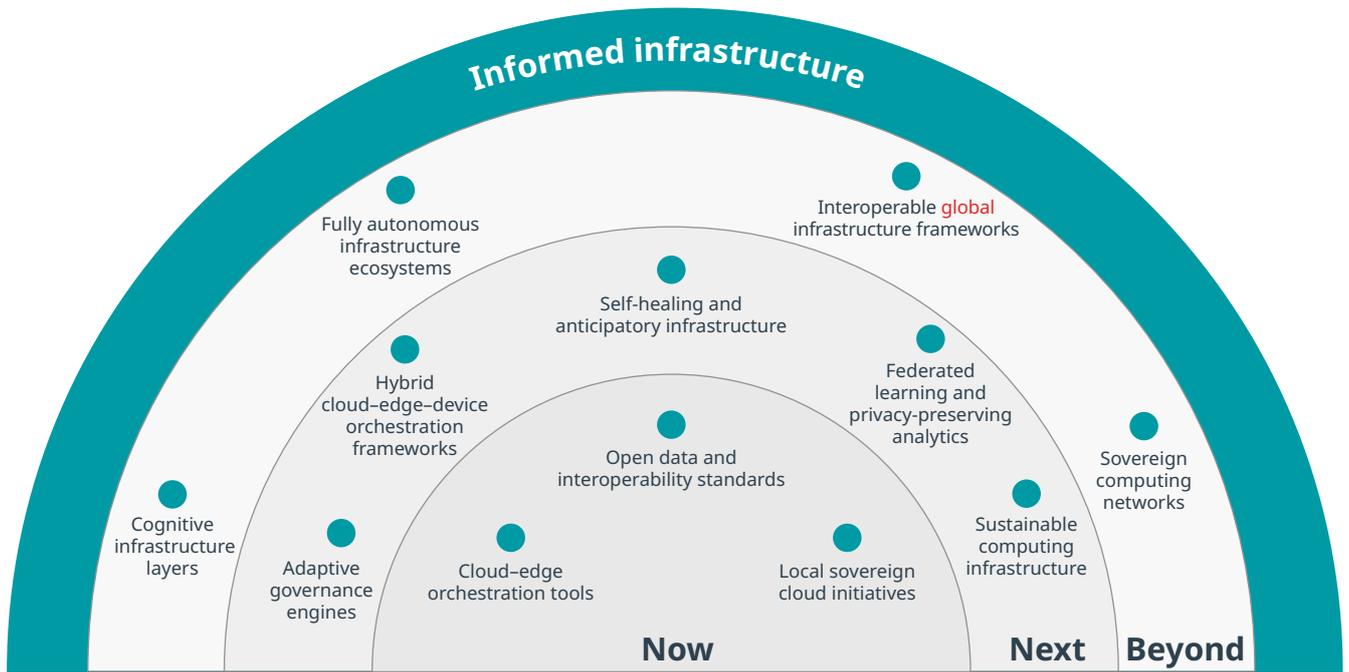


Figure 9: Informed infrastructure: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
AI edge infrastructure	3.65	11.37	25.5%
Global AI infrastructure	56.97	186.27	26.7%
Composable infrastructure market	8.09	32.97	32.4%
Hybrid cloud-edge computing market	0.83	2.71	26.8%

Table 7: Informed infrastructure: Market size and forecast ¹⁶

Trust functions are converging into unified stacks spanning governance, observability, provenance and cybersecurity.

¹⁶ Sources:

Technavio, Global AI Edge Infrastructure Market 2025-2029.

The Business Research Company, Global AI Infrastructure Market Report 2025; Global Composable Infrastructure Market Report 2025.

Enterprise adoption signals

75% to 85%

of GenAI proofs of concept fail because of infrastructure gaps.¹⁷

Approximately

20%

of infrastructures are upgraded for AI workloads.¹⁸



Sovereign cloud and interoperability **accelerate the adoption of audit-ready infrastructure.**

Industry partnership signals

Informed infrastructure is strongest in industries where latency, sovereignty and resilience matter: manufacturing, public sector, energy, telecom, finance and healthcare.

Adoption curve insights

- **2024 to 2026**
The adoption of hybrid cloud-edge architectures, sovereign-cloud zones and telemetry-rich foundations accelerates.
- **2026 to 2028**
Distributed inference, digital twins and self-healing infrastructure scale across industries.
- **2028 to 2030**
Autonomous, sustainability-aware orchestration enables the real-time optimization of computing, energy and operations.
- **2030+**
Intelligent, sovereign infrastructure ecosystems become foundational to both national and enterprise resilience.



^{17,18} Everest Group. Thematic Report. Navigating AI Infrastructure: the Backbone of the AI-Driven Era. September 13, 2024.

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Airbyte	Open source data integration platform (ETL/ELT)* for replicating data from various sources to data warehouses.	2020	Series B	181.5M
Itential	Provides a low-code automation platform specifically for managing and orchestrating network and IT infrastructure.	2014	Growth Stage	25.5M
Liqid	A composable infrastructure company that provides software to dynamically pool and allocate GPUs, CPUs, storage and networking for AI, high-performance computing (HPC), and other data-intensive workloads in data centers.	2013	Post-Series C	160M
OpenDrives	Provider of high-performance storage solutions and data management services for media, entertainment and enterprise.	2011	Series B	31M
Rafay	Offers a Kubernetes operations platform to automate and manage containerized applications across multiple clusters and clouds.	2017	Series B	33M
Systalyze	An AI infrastructure and optimization startup that analyzes and fixes inefficiencies in AI workloads so enterprises can run models faster and more cheaply on their own infrastructure.	2025	—	—

Table 8: Informed infrastructure: Relevant startups ¹⁹

*ELT: extract, load, transform | ETL: extract, transform, load

¹⁹ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate how informed infrastructure enables industries to orchestrate complex physical-digital environments through continuous intelligence, hybrid cloud-edge computing and sovereign digital foundations.

01 Intelligent public infrastructure: Predictive, efficient and sovereign cities

Cities and public agencies adopt predictive maintenance for public assets, smart traffic orchestration, and water- and waste-optimization systems to improve service reliability and resource efficiency.

Cloud-edge orchestration and digital twins allow infrastructure operators to detect failures early, optimize flows in real time, and test interventions virtually before deployment. These systems increasingly run on sovereign-cloud foundations, enabling compliance, data residency and secure cross-agency collaboration.

02 Adaptive industrial operations: Digital twins, edge AI and self-healing systems

Manufacturing, logistics and energy companies deploy industrial digital twins and edge AI inference in hybrid cloud-edge environments to achieve low-latency decision-making and resilient operations.

AI-powered observability and self-healing infrastructure enable the continuous optimization of production lines, predictive maintenance for equipment and autonomous recovery from disruptions.

The shift to distributed computing fabrics, real-time telemetry and anticipatory infrastructure reduces downtime and improves performance across industrial networks.

03 Sustainable, distributed energy and data ecosystems

Energy providers and large enterprises advance dynamic energy management, net-zero data center strategies, and sovereign digital corridors that coordinate computing, storage and energy production and distribution across regions.

Digital twins model emissions and energy behavior at grid scale, while sustainability-aware scheduling routes workloads based on energy availability and carbon intensity. Hybrid cloud-edge orchestration supports local energy optimization, grid balancing and the integration of distributed renewable sources.

A strong investment momentum in sustainable computing infrastructure and energy-optimized operations is driven by regulatory pressure and the steady growth of energy-aware infrastructure technologies.

Future scenarios

Auto-remediation becomes the default in cloud operations

Most cloud incidents are resolved without human intervention. Operations teams shift from fixing problems to validating policies.



Primary uncertainties in this scenario

- Will companies accept shrinking human oversight?
- Will regulators require human override paths for critical workloads?

Digital-twin networks replace traditional monitoring

Enterprises adopt real-time digital twins of infrastructure as the operational “source of truth,” simulating changes before implementing them.



Primary uncertainties in this scenario

- Can digital twins achieve sufficient fidelity to remain trusted at scale?
- Will simulation latency or cost become a limiting factor?

Edge-first architectures overtake cloud-first thinking

Computing orchestrates from edge devices inward; cloud serves as coordination rather than central computing.



Primary uncertainties in this scenario

- Will orchestration intelligence mature enough for edge-optimized silicon to replace centralized GPU computing for most real-time decisions?
- Will organizations embrace distributed orchestration or fear loss of control?



Strategic takeaways



Treat infrastructure as a strategic intelligence layer, not a technical utility

Organizations must evolve from static IT foundations to AI-optimized, telemetry-rich, hybrid cloud-edge architectures. This means proactively upgrading networks, storage and computing for real-time inference, sovereignty and sustainability.



Build sovereignty and interoperability into architecture decisions now

With sovereign-cloud adoption accelerating and interoperability standards becoming essential, infrastructure choices must ensure governance, jurisdictional control and vendor modularity. Investing in open standards, composability and specialized silicon reduces dependency risk and prepares organizations for regulatory shifts.



Operationalize sustainability through intelligence, not reporting

Energy-aware scheduling and digital twins transform environmental, social and governance (ESG) considerations from being a compliance burden to real-time operational optimization. Organizations that integrate sustainability into infrastructure operations — not just dashboards — unlock measurable reductions in cost, energy use and emissions.



- If our infrastructure were fully self-optimizing, sovereign-aligned and energy-aware, how would we redesign our operating model to take advantage of it?
- Do we understand which decisions should be automated by intelligent infrastructure and which must remain human-led?
- How resilient is our infrastructure strategy to shocks in energy supply, geopolitics or regulatory sovereignty?

5 Sovereign silicon ecosystems

What it is and why it matters now

The hardware foundation of mass intelligence

Geopolitical tensions, supply chain fragility, AI-driven computing demand and rising national security concerns have made semiconductors a defining strategic asset. Control over silicon now equates to control over computing capacity, data flows, intelligence and geopolitical influence. Nations and industries are therefore accelerating efforts to develop end-to-end semiconductor ecosystems, from chip design and photonics to fabrication, packaging and recycling.

This shift matters now because traditional global supply models can no longer guarantee resilience, cost stability or technological competitiveness.

At the same time, rapid advances in AI, high-performance computing and next-generation architectures require specialized, heterogeneous computing that legacy supply chains cannot deliver alone. Sovereign silicon ecosystems offer a path to innovation leadership, technological independence, sustainable manufacturing and collaborative advantage.

They also anchor the hardware foundation of mass intelligence, supplying the specialized, resilient and high-performance computing required for autonomous, data-rich and globally distributed systems.

Underlying concepts

Semiconductors as enablers of AI and HPC

Chips power AI and high-performance computing, defining competitiveness in data-driven innovation

Industrial and technological autonomy

Domestic semiconductor capabilities reduce external dependencies and strengthen national innovation capacity.

Data sovereignty and digital infrastructure control

Owning the silicon layer provides trusted foundations for secure and sovereign digital ecosystems.

Innovation ecosystems and collaborative advantage

Public-private-academic partnerships accelerate chip innovation aligned with national priorities.

Sustainable and ethical supply chains

Responsible sourcing, recycling and energy-efficient manufacturing influence future semiconductor competitiveness.

Technology-driven competitive advantage

Leadership in advanced manufacturing and the creation of intellectual property (IP) drives global influence.

AI-accelerated design and manufacturing

AI shortens chip-development cycles and improves yields, enhancing resilience and agility.

Photonics as the next frontier of semiconductor sovereignty

Optical integration enables ultrafast, energy-efficient computing for next-generation applications.

Resilient and adaptive global networks

Flexible, collaborative semiconductor ecosystems balance sovereignty with global innovation flows.

Sovereignty in silicon is not a single capability but a system of strategic levers spanning design, manufacturing, materials, integration, knowledge and energy.

Sovereignty levers



Figure 10: Sovereignty levers



Technology landscape

Sovereign silicon ecosystems are enabled by a rapidly advancing stack of fabrication technologies, heterogeneous computing architectures, trusted hardware and energy-efficient acceleration. These technologies shift silicon from being a globalized commodity to becoming a strategically governed infrastructure — one that delivers secure, high-performance and sovereign computing capacity for AI, HPC and critical national systems.

● **Now:** Mature capabilities forming today's sovereign computing infrastructure

Today's deployments rely on advanced semiconductor fabrication, trusted foundry models and regional semiconductor hubs, supported by massive public investment.

These capabilities strengthen manufacturing sovereignty and reduce exposure to concentrated global supply chains.

Performance and flexibility are enabled by chiplet architectures and AI accelerator chips, a segment expected to dominate data center and sovereign-cloud deployments as AI demand increases exponentially.

Sovereign transparency and trust are supported by secure hardware roots of trust, confidential computing and semiconductor traceability models, while sustainable chip manufacturing and advanced packaging deliver performance-per-watt gains, supporting both resilience and energy-efficiency objectives.

● **Next:** Scaling technologies enabling sovereign AI and distributed computing

The next wave introduces deeper intelligence, security and sustainability across the silicon lifecycle.

AI-assisted chip design and AI-accelerated fabrication compress development cycles and increase yield. These outcomes are critical as design complexity and national demand for sovereign AI computing accelerate. Edge AI and scaling inference extend sovereign computing to industrial sites, mobility systems and national infrastructure, aligning silicon with autonomy and real-time operational needs.

Sustainability becomes systemic through carbon-aware computing and carbon-neutral manufacturing, enabling nations to balance computing growth with environmental commitments.

Security hardening intensifies: Trusted AI-on-silicon platforms, quantum-safe encryption and post-quantum cryptography (PQC) protect national workloads against advanced and emerging threats, supporting compliance with sovereignty and data-residency mandates.

At the architectural layer, hybrid photonic–electronic integration begins to unlock high-bandwidth, low-energy interconnects for HPC and sovereign-cloud infrastructure.

● **Beyond:** Frontier technologies creating autonomous, high-performance sovereign computing

Longer-term development points to silicon architectures capable of supporting autonomous national computing environments.

Full-scale silicon photonics and optical computing architectures offer ultrahigh throughput and radically improved energy efficiency for simulation, climate modeling and next-generation AI workloads.

Meanwhile, bio-inspired and advanced neuromorphic chips introduce event-driven, spiking and analog computation models that enable real-time sensing, autonomous robotics and energy-constrained national edge systems at unprecedented scale.

Together, these frontier technologies lay the groundwork for sovereign computing that is not only secure and self-determined but also adaptive, ultra-efficient and capable of powering future national intelligence infrastructures.

Technology radar

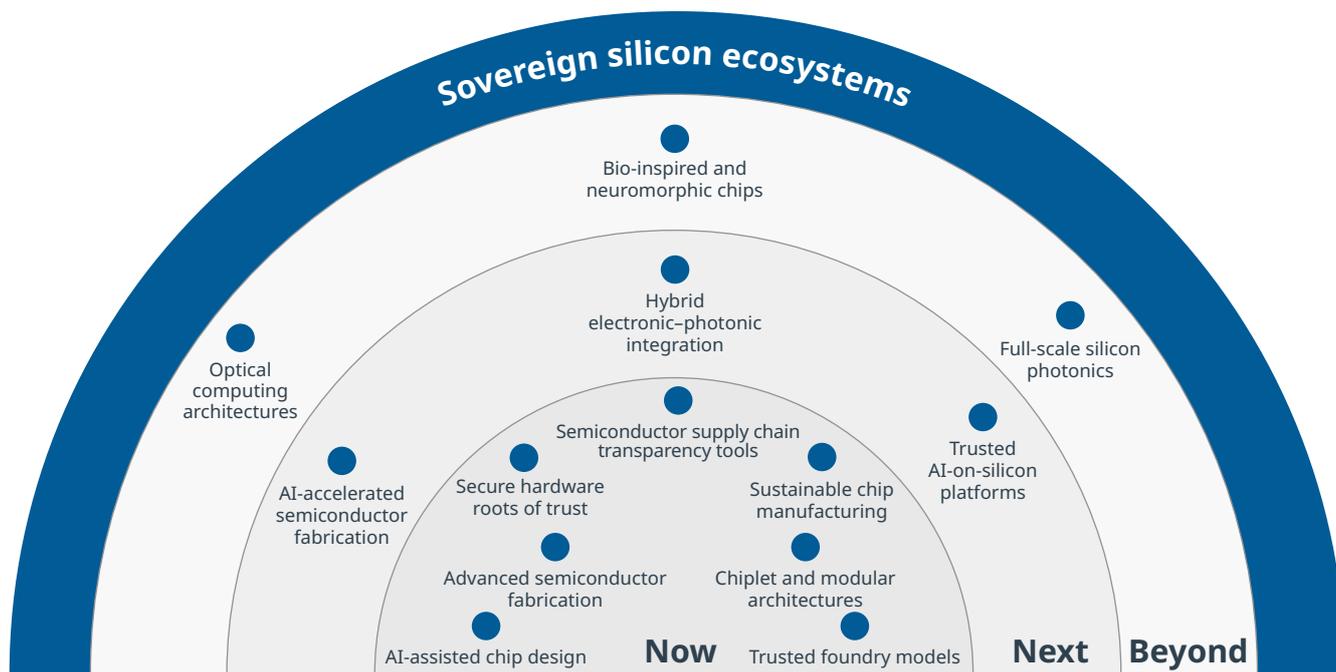


Figure 11: Sovereign silicon ecosystems: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
Global semiconductor market	540.7 *	892.6 *	10.6% *
Advanced packaging	42.0	76.0	10.6%
Global sovereign cloud	124.1 *	377.7 *	24.6%
AI chip market	120.0	500.0	35.0%

Table 9: Sovereign silicon ecosystems: Market size and forecast ²⁰

*Internal estimates

AI-driven computing demand is fueling sovereign semiconductor investment at unprecedented scale.

²⁰ Sources:

World Semiconductor Trade Statistics. Press release. [Global Semiconductor Market Approaches \\$1T in 2026](#), December 2, 2025

Fortune Business Insights. [Semiconductor market size, share and industry analysis, 2026-2034](#), February 9, 2026.

Fortune Business Insights. [Sovereign cloud market size, share and industry analysis, 2026-2034](#), February 9, 2026.

Everest Group. Thematic Report. [From Silicon to Strategy: Mapping the AI Chip Value Chain](#), September 2, 2025.

Enterprise adoption signals

Servers with an embedded accelerator are the preferred infrastructure for AI platforms, accounting for

70%

of total spending on server AI infrastructure.²¹

The CHIPS and Science Act allocates roughly

\$52.7 billion

to strengthen US semiconductor manufacturing and research through incentives, R&D funding and workforce programs.²²



The EU launched the European Chips Act to strengthen semiconductor production capacity and related R&D in member states through public funding incentives and regulatory support, as part of the EU's industrial policy.

Industry partnership signals

The need for sovereign silicon is highest in the defense, telecom, automotive and industrial automation sectors, all of which require trusted, high-performance computing.

Adoption curve insights

2025 to 2027

Sovereignty gaps become visible as demand for trusted, high-performance computing outpaces existing supply chains.

2027 to 2029

Select nations and enterprises mature their sovereign silicon stacks through domestic design, fabrication and packaging.

2029 to 2032

Heterogeneous, energy-efficient architectures — chiplets, photonics, AI accelerators — become mainstream computing foundations.

2032+

Sovereign computing grids and advanced silicon ecosystems have an impact on global alliances and digital-economy power structures.



²¹ IDC. Press release. Artificial Intelligence Infrastructure Spending to Surpass the \$200Bn USD Mark in the Next 5 years. According to IDC, February 18, 2025.

²² National Institute of Standards and Technology (NIST). Federal Programs Supporting the U.S. Semiconductor Supply Chain and Workforce. March 18, 2024.

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Axelera AI	Developer of ultralow-power and high-performance AI hardware for edge applications.	2021	Series B/Grant	203M
Lightmatter	Designs and builds photonic computer chips that use light instead of electricity for faster and more efficient AI/ML computation.	2017	Series D	850M
Silicon Catalyst	An incubator/accelerator focused exclusively on enabling semiconductor hardware startups.	2015	N/A (Accelerator)	—
Xscape Photonics	Develops advanced silicon photonics technology for faster and more energy-efficient data communication and computing.	2022	Series A	57M

Table 10: Sovereign silicon ecosystems: Relevant startups ²³

²³ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate how sovereign silicon ecosystems translate from industrial policy and supply chain control into concrete national capabilities for critical infrastructure, industry and global systems.

01 Sovereign AI computing grids for critical industries

To support regulated sectors such as healthcare, finance, transportation and public services, nations deploy AI sovereignty platforms — domestic AI accelerator clusters (GPUs and other purpose-built accelerators) built on trusted foundry chips and sovereign-cloud infrastructure.

Energy-efficient AI accelerators and photonics-enabled green computing allow industries to run high-performance models with significantly lower power consumption, supporting sustainability mandates while reducing dependency on foreign computing providers.

This sovereign-computing grid functions as shared critical infrastructure for the digital economy, enabling privacy-preserving analytics, national digital twins, autonomous mobility optimization and secure collaboration across sectors.

02 Secure industrial autonomy and national edge-control systems

The manufacturing, energy, mobility and logistics sectors adopt secure AI inference at the edge using locally produced control electronics and trusted industrial controllers.

This enables operational autonomy for factories, grids, transportation hubs, ports and defense logistics — environments where latency, safety and geopolitical risk prohibit reliance on external silicon or cloud services.

PQC-ready modules and sovereign-identity hardware secure every device, while edge inference allows for real-time anomaly detection, predictive maintenance and autonomous decision-making in multiple industries.

The result is resilient national infrastructure, and safety-critical decisions remain within sovereign boundaries.

03 Strategic silicon platforms for space, communications and global alliances

Telecom operators, aerospace agencies and defense sectors adopt secure communication networks, sovereign AI accelerators and radiation-tolerant chips to power national satellite and space systems.

These sovereign silicon platforms support encrypted satellite links, disaster-response coordination, climate monitoring, maritime safety and critical cross-border communications.

Because the chips and control systems are domestically governed, nations gain leverage in innovation diplomacy — forming strategic semiconductor alliances, influencing standards, and using sovereign silicon as an asset in trade negotiations and international partnerships.

Industries ranging from agriculture and insurance to logistics and mining rely on these sovereign space and communication systems for precision services and resilient connectivity.

Future scenarios

Chiplet marketplaces emerge from cloud providers

Cloud vendors begin offering modular chiplet configurations so organizations can tailor accelerators to their workloads.



Primary uncertainties in this scenario

- Will chiplet standards mature fast enough for broad compatibility?
- Will IP owners allow meaningful configurability?

Regional computing zones change global cloud footprints

Nations create regional sovereign computing hubs with mandatory domestic supply chains, forcing clouds to distribute AI workloads geopolitically.



Primary uncertainties in this scenario

- Will national investments achieve competitive performance?
- Can data or model portability work reliably across sovereign zones?

Neuromorphic, quantum and photonic systems reset the computing paradigm

When traditional scaling plateaus, enterprises adopt neuromorphic and photonic accelerators for low-energy inference and planning workloads.



Primary uncertainties in this scenario

- Will software ecosystems mature quickly enough to support these architectures?
- Can heterogeneous hardware orchestration become seamless?



Strategic takeaways



Treat silicon as national and enterprise-critical infrastructure

Organizations should assume that control over chips, accelerators and secure computing capacity will determine their ability to operate, innovate, and comply with sovereignty requirements. Prioritizing partnerships with trusted foundries, domestic packaging ecosystems and sovereign-cloud providers becomes essential to securing long-term resilience.



Build a heterogeneous, energy-efficient computing strategy

Chiplet architectures, energy-efficient AI accelerators and photonics-enabled systems allow companies to scale AI workloads sustainably and cost-effectively. Enterprises should architect for modularity, mix-and-match accelerators and edge inference to reduce their reliance on single-vendor or single-region supply chains.



Embed hardware-rooted trust into every AI and infrastructure deployment

Secure identity hardware, PQC-ready modules and trusted industrial controllers must become the baseline for any AI or critical-infrastructure application. This facilitates data integrity, operational safety and compliance across distributed systems, from factories and energy grids to transportation and national services.



- If our organization lost access to our current chip supply or cloud accelerators for six months, how much of our AI, operational technology and digital infrastructure could we continue to run? What would need to be sovereign for us to remain resilient?
- Are we designing our AI and infrastructure workloads to be portable across heterogeneous, sovereign computing environments?
- Which parts of our silicon ecosystem must be controlled directly? And where do strategic partnerships create more resilience than full ownership?

6 From illusory efficiency to sufficiency

What it is and why it matters now

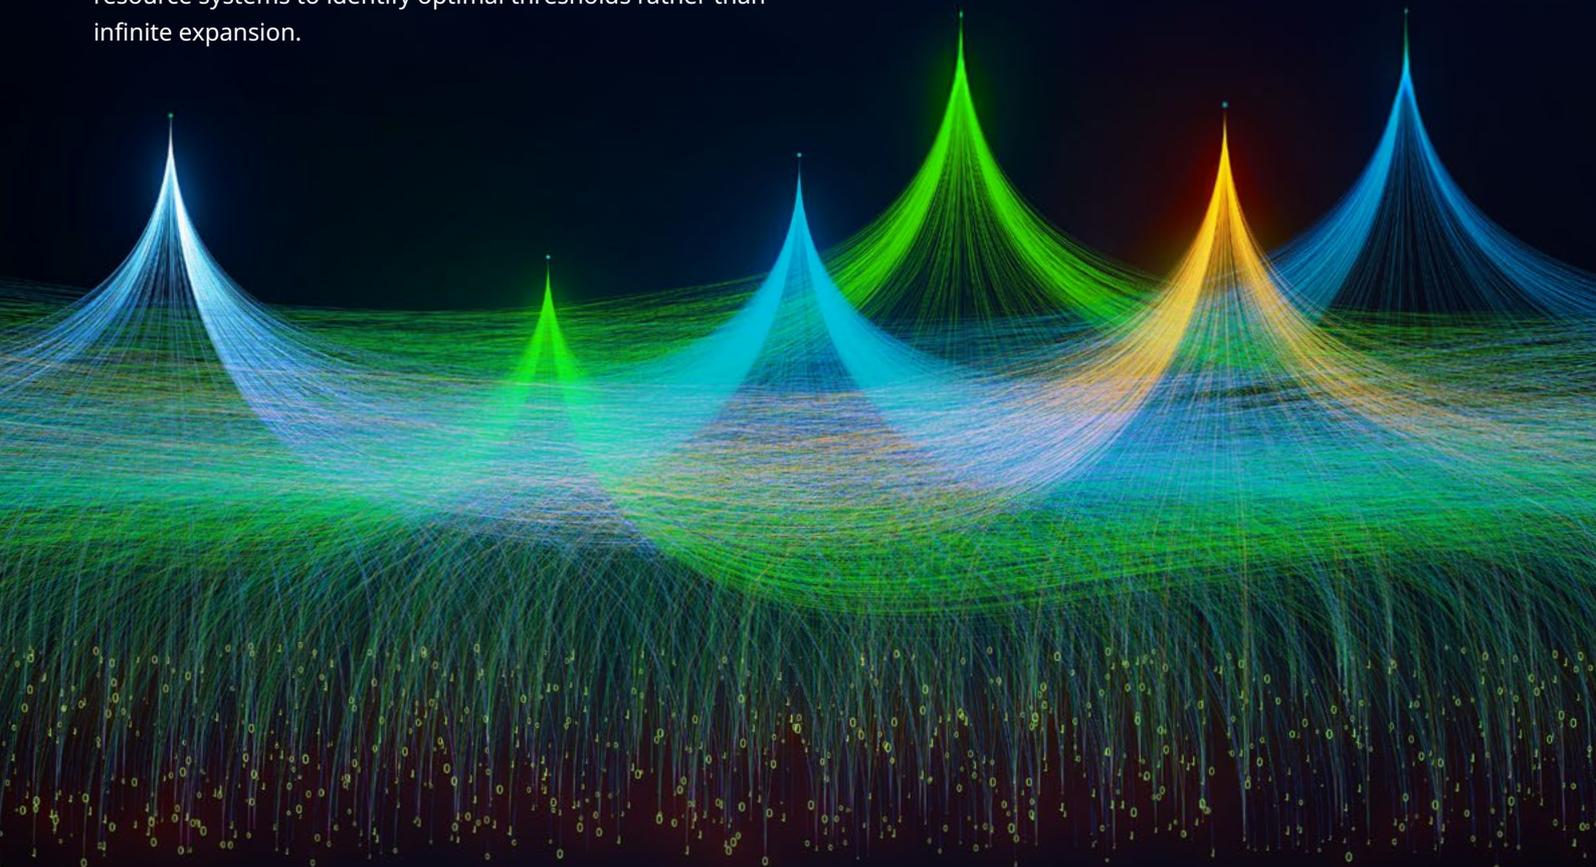
Reframing how mass intelligence evolves

After decades of optimizing for speed, scale and cost, organizations are confronting the limits of efficiency-driven growth. Climate instability, energy scarcity, regulatory pressure and social expectations require a fundamentally different model — one that prioritizes resilience, long-term adequacy and ecological alignment. Sufficiency matters today because efficiency alone can no longer guarantee stability or credibility; in some cases, it accelerates systemic risk by maximizing short-term gains at the expense of planetary boundaries.

The sufficiency paradigm reframes technology as a tool for intelligent moderation, using AI, digital twins and integrated resource systems to identify optimal thresholds rather than infinite expansion.

This approach unlocks shared value in the form of lower environmental impact, greater community resilience, better working conditions and more trustworthy business practices. Sufficiency signals a cultural and strategic shift to where progress is defined not by “more” but by purposeful, responsible and sustainable growth.

Sufficiency also reframes how mass intelligence evolves, ensuring that intelligent systems operate within ecological and societal boundaries, and channel their capabilities to long-term resilience rather than unchecked expansion.



Underlying concepts

Sufficiency as a strategic mindset

Organizations shift from maximizing output to optimizing for long-term adequacy within ecological limits, and thus future-proof their business models.

Systemic resource stewardship

The integrated management of energy, water and materials minimizes total resource footprints across systems.

Embedded environmental accountability

Lifecycle impacts — carbon, biodiversity, water — are assessed and mitigated at every planning, design and operational stage.

Resilience over optimization

Performance prioritizes adaptability and robustness rather than minimal cost.

Transparency and credibility loops

Digital tools (IoT, lifecycle assessment (LCA), blockchain) ensure traceable, audit-ready sustainability performance.

Human and ecological co-benefits

Sufficiency improves working conditions, community resilience and ecosystem health.

Adaptive governance and shared accountability

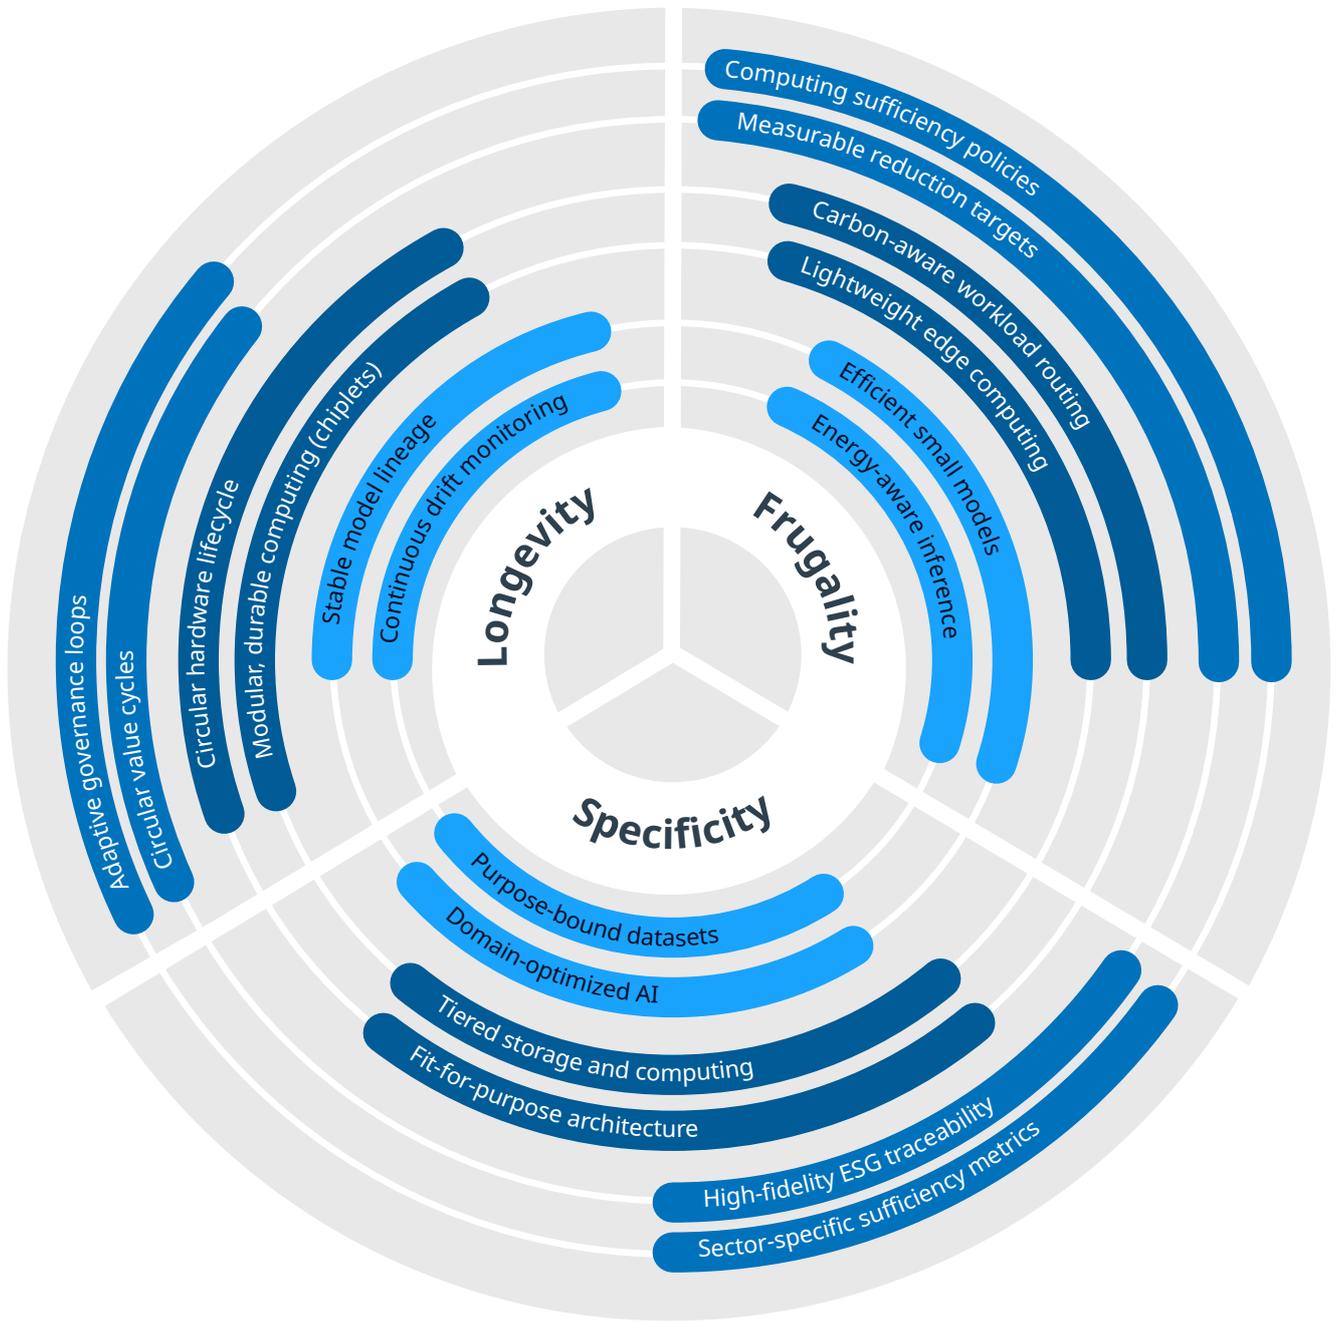
Co-created standards and metrics align operational choices with societal goals.

Technological enablers for responsible scale

AI, automation and digital twins identify optimal sufficiency thresholds and link local efficiencies to systemic transitions.

The sufficiency circle maps the core dimensions of technological sufficiency, showing how AI, digital infrastructure and governance must evolve to balance performance with responsible, long-term resource use.

Technological sufficiency circle



Governance and societal sufficiency
(The systemic alignment layer)

Digital infrastructure and systems sufficiency
(The architecture core)

AI and data sufficiency
(The intelligence core)

Figure 12: Technological sufficiency circle

Technology landscape

The technology landscape of the macrotrend “from illusory efficiency to sufficiency” is defined by capabilities that shift organizations from incremental optimization to verifiable, measurable reduction.

● Now: Operational reduction systems

Organizations are deploying technologies that deliver immediate, quantifiable resource savings.

IoT-based resource monitoring and AI-enabled energy-management systems provide continuous visibility of consumption patterns, enabling reductions in energy, water and waste within just a few years. Circular manufacturing systems, sustainable design platforms and sustainability-by-design practices embed durability, repairability and minimal-input design into core product and service development.

At the infrastructure level, smart buildings and smart-city frameworks optimize resource use in heating, cooling, lighting, mobility and utilities, while renewable-energy management systems orchestrate low-carbon energy flows. AI carbon-accounting platforms and circular IT asset management systems provide the audit-grade data required to comply with requirements and standards such as those of the Corporate Sustainability Reporting Directive (CSRD), Corporate Sustainability Due Diligence Directive (CSDDD) and International Sustainability Standards Board (ISSB), marking a wider shift from reporting to evidence-based reduction.

● Next: Predictive circularity and low-carbon computing

As sufficiency matures, enterprises adopt systems that enable predictive, cross-value-chain resource optimization.

Carbon-aware computing and net-zero data centers reduce the footprint of digital operations by aligning workloads with low-carbon energy availability. All-photonic networks emerge as a key enabler, mitigating latency and energy loss to support dynamic, low-carbon workload distribution across data centers and regions. Digital twins for sustainability model energy, materials and emissions flows before changes occur, helping organizations anticipate reduction outcomes rather than react to them.

Industrial ecosystems advance through carbon-neutral chip manufacturing, AI-augmented circular logistics and low-carbon material innovation, extending sufficiency into supply chains, manufacturing and semiconductor infrastructure. These capabilities reflect the report’s forecast that circular-economy and sustainability software markets will grow steadily through 2030.

● Beyond: System-level sufficiency infrastructure

Looking further ahead, sufficiency transitions from enterprise practice to societal infrastructure.

Planetary digital twins integrate climate, infrastructure and economic systems to model long-term sufficiency pathways within planetary boundaries. AI sufficiency policy simulators allow governments and cities to test demand-reduction strategies, circular procurement and regulatory interventions before implementation. At the energy-system level, next-generation low-carbon sources — most notably nuclear fusion — are positioned as potential foundations for stable, abundant baseload power compatible with sufficiency. While still precommercial, AI- and quantum-enabled advances in simulation, plasma control and material informatics are accelerating fusion’s transition from scientific breakthrough to deployable energy infrastructure.

Bio-intelligent materials and decentralized regenerative economies signal a shift to regenerative value creation, where materials, energy and production cycles are inherently circular. These long-horizon innovations align with the report’s emphasis on adaptive governance, human-ecological co-benefits and systemic resilience as the ultimate markers of sufficiency.

Technology radar

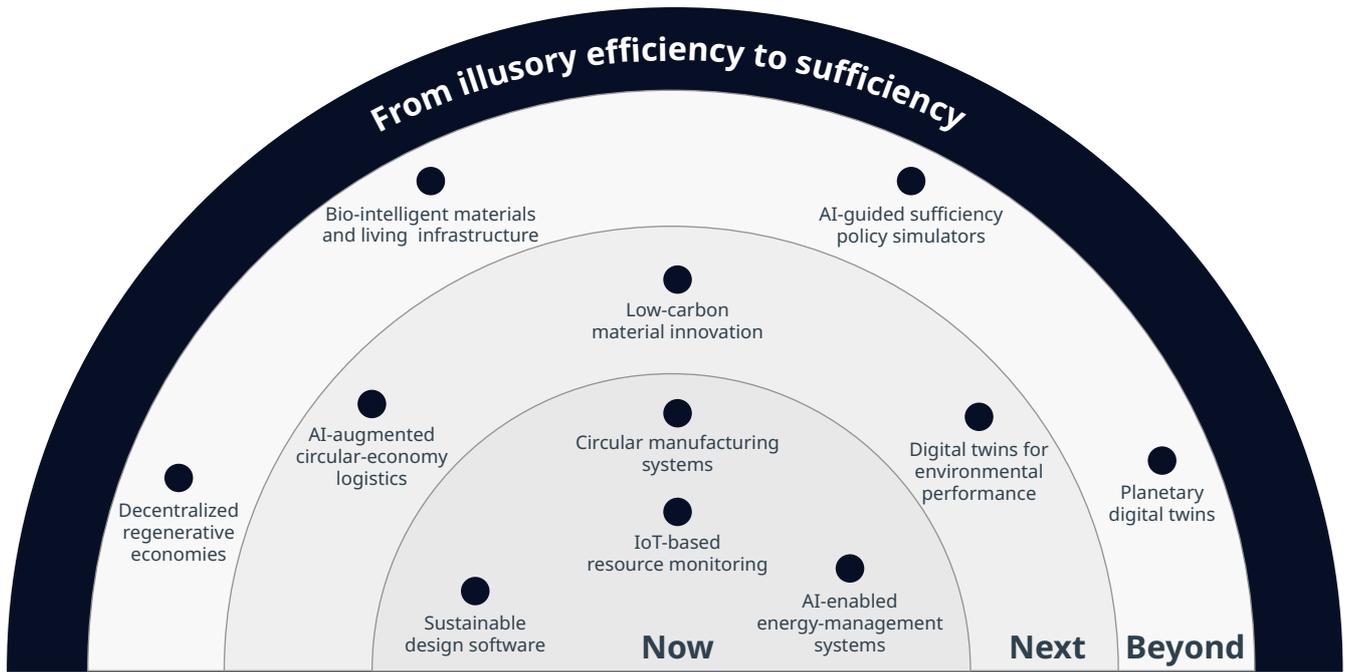


Figure 13: From illusory efficiency to sufficiency: Technology radar

Signals and ecosystem

Market signals

Segment	Market size 2024 (USD B)	Market size 2029 (USD B)	CAGR (2024 to 2029)
Green technology and sustainability	23.10	65.30 *	23.1%
Circular economy	465.20 *	794.40 *	11.3%
Sustainability consulting services	36.40 *	114.20 *	25.7%
Sustainability/ESG management software	3.21 *	6.88 *	16.7%

Table 11: From illusory efficiency to sufficiency: Market size and forecast ²⁴

*Internal estimates

Sufficiency is evolving from compliance to measurable operational reduction.

²⁴ Sources:

Grand View Research, [Market Analysis Report: Green Technology & Sustainability Market \(2025–2030\)](#).

The Business Research Company, [Circular Economy Market Report 2026](#), March 2026.

Mordor Intelligence, [Sustainability Consulting Services Market Size & Share Analysis – Growth Trends and Forecast \(2026–2031\)](#), January 2026.

Grand View Research, [Market Analysis Report: Sustainability Management Software Market \(2024–2030\)](#).

Enterprise adoption signals

By mid-2025, over

1,400

corporates had set net-zero targets, representing **38%** of all corporates with such targets. ²⁵

By 2030, more than

65%

of global enterprises will use agentic-AI-driven ESG software to support sustainable sourcing. ²⁶



Regulatory drivers (CSRD, ISSB, US Securities and Exchange Commission (SEC)) **accelerate verifiable ESG data flows and traceability.**

Industry partnership signals

The manufacturing, finance, retail, healthcare, logistics and energy sectors lead the adoption of traceability, lifecycle intelligence and verified reduction.

Adoption curve insights

- 2024 to 2026**
ESG data systems, traceability platforms and lifecycle intelligence gain rapid adoption, driven by regulation.
- 2026 to 2028**
Operational sufficiency scales through circular design, resource-intelligent processes and digital-twin optimization.
- 2028 to 2030**
Verified reductions in energy, materials and emissions become standard performance metrics.
- 2030+**
Sufficiency becomes a structural operating model, guiding investment, innovation and long-term resilience.



^{25, 26} Science Based Targets, [Data and Trends - Science Based Targets Initiative](#).

Relevant startups

Startup name	Overview	Founded in	Funding stage	Total funding (USD)
Akselos	Provides a simulation technology platform using digital twins for industrial assets (ships; energy infrastructure).	2012	Series A/B	29.5M
CleanHub	Operates a platform that helps brands fund and manage the recovery of plastic waste from the environment.	2020	Seed/Venture	11.9M
Felt	A collaborative, web-based mapping and geospatial data visualization tool.	2021	Venture/ Series unknown	34.5M
Neara	Software platform using digital twins and AI to model, simulate and manage complex utility and infrastructure networks.	2016	Series C	60.8M
Rheaply	Enterprise platform for internal resource exchange and asset management, promoting circularity and sustainability.	2015	Venture/ Series unknown	32.8M
TwinUp	Builds real-time, geospatial digital twins for cities, infrastructure and industry, enabling decision-makers to simulate, monitor and optimize physical systems virtually.	2022	Unknown	7M

Table 12: From illusory efficiency to sufficiency: Relevant startups ²⁷

²⁷ The information presented here is based on data obtained from Crunchbase (<https://www.crunchbase.com/>). Accuracy and completeness depend on the source and may change over time.

Use cases and applications

The following examples illustrate how sufficiency is becoming a practical operating model across industries, enabled by digital intelligence, circular value chains and verifiable environmental data.

01 Resource-intelligent infrastructure — microgrid sufficiency and data-driven buildings

Enterprises and municipalities are deploying microgrid sufficiency models that optimize local energy generation, storage and consumption.

AI-enabled energy-management systems balance loads between solar, wind, hydrogen, battery assets and grid inputs, enabling communities to operate within self-defined sufficiency thresholds rather than maximizing throughput.

In parallel, data-driven building operations leverage IoT monitoring, digital twins and predictive analytics to reduce heating, cooling and lighting demand.

These capabilities enable significant campus-level energy reductions while creating audit-ready evidence for CSRD and ISSB reporting.

02 Circular production and low-waste supply chains — AI for waste minimization

Manufacturers increasingly use circular manufacturing loops supported by lifecycle analytics, material-traceability systems and AI-driven waste minimization.

Digital twins model component lifecycles to determine whether items should be repaired, remanufactured or recycled, helping companies shift from end-of-life disposal to regenerative cycles.

Logistics networks integrate AI-augmented circular routing to optimize loads, orchestrate reverse logistics, and reduce material loss in warehouses and transport flows. These systems transform linear supply chains into low-resource circular ecosystems.

03 Planetary-aware decision-making — precision agriculture and boundary dashboards

Precision agriculture uses AI-driven soil analytics, multispectral imaging and micro-irrigation optimization to increase yield while reducing water, fertilizer and fuel consumption — aligning agricultural productivity with sufficiency rather than maximization.

These systems put emphasis on systemic resource stewardship and help stabilize input use under tightening environmental constraints.

On a broader scale, planetary boundary dashboards integrate satellite data, emissions inventories, water-use metrics, and circularity indicators to guide decision-makers toward “enough”.

This creates a global sufficiency index that signals when systems are approaching ecological limits. Governments and enterprises use these dashboards to align policy, procurement and investment decisions with verified environmental thresholds.

Future scenarios

Carbon-aware workload routing becomes default

AI and cloud workloads move automatically to regions with lower carbon intensity, even at the cost of latency or price.



Primary uncertainties in this scenario

- Will customers tolerate latency trade-offs for greener computing?
- Will energy and carbon APIs become standardized across regions?

Computing-sufficiency policies replace cost optimization

Enterprises impose caps on computing consumption per model, per team or per customer, shifting innovation from scaling to efficiency.



Primary uncertainties in this scenario

- Can new hardware architectures reduce energy per compute quickly enough to avoid hard caps?
- Will competitive pressures push companies to violate sufficiency policies?

Model size reversal becomes industry standard

Hyperspecialized, domain-specific small language models outperform massive general models for most enterprise needs, creating a "post-scale" AI economy.



Primary uncertainties in this scenario

- Will small-model tooling and lifecycle management reach industrial maturity?
- How quickly can organizations migrate from generalist LLMs to vertical SLMs?



Strategic takeaways



Shift from optimization to absolute reduction

Efficiency is no longer enough. Organizations must redesign operations, products and infrastructures to deliver verifiable reductions in total energy, materials and emissions, not just incremental improvements. This requires embedding sufficiency-by-design into research, product development, procurement, production, service and decommissioning.



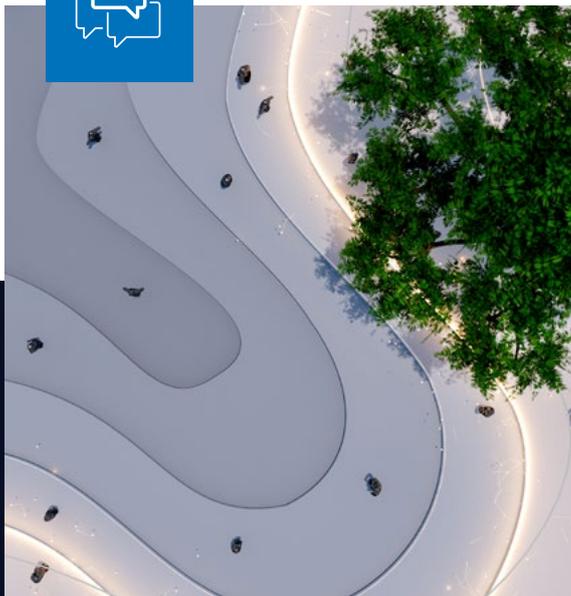
Build a digital backbone for traceability and accountability

Regulation now demands audit-ready, high-fidelity environmental data. Enterprises should prioritize ESG data pipelines, carbon-accounting platforms, lifecycle analytics and circular-asset management systems to create a trusted foundation for compliance and reduction decisions. Digital evidence — not estimates — will determine both regulatory readiness and capital-market credibility.



Operationalize circularity across value chains

Sufficiency becomes scalable when circular flows — repair, reuse, remanufacturing, recovery — are operationalized through AI, digital twins and traceability. Industrial systems that adopt circular loops and waste-minimization analytics already achieve significant reductions in material losses. Enterprises should embed circularity into supply chain design, logistics, IT asset lifecycles and manufacturing processes to unlock near-term efficiency and long-term resilience.



- If efficiency is no longer the path to climate alignment, which part of our operating model must fundamentally change — not optimized, but reinvented — to achieve “enough” rather than “more”?
- Which stakeholders define sufficiency for our organization, and how are those boundaries governed over time?
- Are we prepared to redesign products, services and incentives around durability, reuse and restraint — even if this challenges short-term revenue models?

Technology in context: Foresight across industries

The six macrotrends gain their full meaning when viewed through the lens of industry transformation. Each sector has a different regulatory environment, customer expectations, operational realities and competitive pressures, and thus experiences mass intelligence differently.

The perspectives presented here illustrate how technological convergence translates into industry transformation, and reveal how the same six macrotrends manifest differently in value chains, regulatory landscapes and operating models.

From factories and hospitals to banks, mobility ecosystems, telecom networks and insurers, the interplay between autonomy, emotion-aware systems, trustworthy intelligence, sovereign infrastructure and sufficiency-driven design influences not only competitiveness but also the resilience of entire industries.

The heatmap below highlights where the six macrotrends exert the strongest impact across sectors, providing a cross-industry lens before diving into each perspective.

Trend relevance heatmap

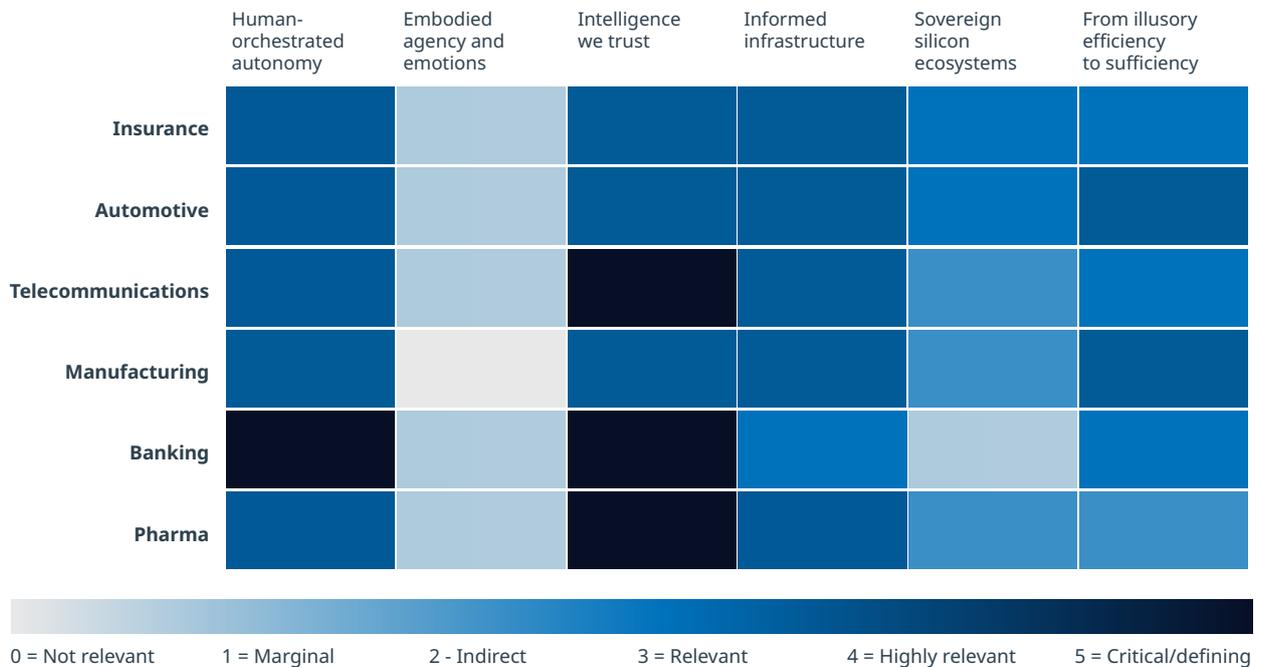


Figure 14: Trend relevance heatmap

Industry perspectives

Insurance

Insurance enters 2026 facing systemic risk inflation, embedded distribution pressures, data-sovereignty requirements and the operational disruption brought about by agentic AI, automation and IoT.

Climate-driven natural catastrophe (NatCat) severity, cyber accumulation, longevity risk and geopolitical fragmentation widen the protection gap as reinsurance capacity tightens and capital becomes more expensive.

At the same time, customers expect embedded, preventive, emotionally intelligent and hyperpersonalized insurance experiences — often delivered through third-party platforms. Insurtech competitors accelerate this shift by redesigning pricing, underwriting and claims around real-time data and continuous risk sensing.

As a result, the industry is moving from reactive indemnification to predictive, preventive and data-driven resilience, supported by intelligent agents, digital twins and partner ecosystems.

Main industry challenges



Systemic risk inflation and portfolio resilience

Climate, cyber, longevity and geopolitical risks outpace traditional models, demanding richer data, digital twins and more efficient reinsurance strategies.



Embedded distribution, platform dynamics and customer trust

Access shifts to retail, mobility and financial platforms, requiring API-native integration, transparent economics and emotion-aware engagement.



AI-driven underwriting, claims and operations

Automation boosts speed and accuracy but depends on explainability, bias control, fairness and regulatory trust.



Cybersecurity, data sovereignty and compliance by design

Correlated cyber events, privacy rules and localization mandates drive modernization based on modular, secure-by-design architectures.



Prevention, circular models and sustainability integration

IoT, telematics, wearables and behavioral data shift insurers to continuous prevention and outcomes-based pricing.

Trend relevance on the insurance value chain

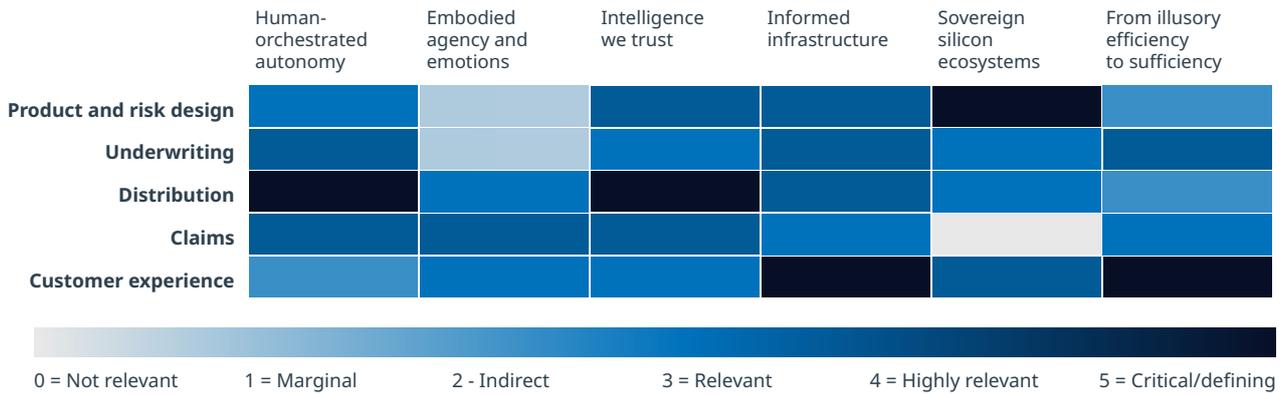


Figure 15: Insurance value chain: Trend relevance heatmap

Transformation drivers

Agentic insurance operations and human-guided automation

AI agents orchestrate underwriting, claims servicing and back-office workflows across multiple steps, increasing speed and consistency while maintaining human accountability, escalation and regulatory control.

High-resolution risk intelligence and prevention

Geospatial analytics, digital twins, IoT, telematics, wearables and biosignals provide near-real-time insight into physical, cyber and behavioral risk, enabling more accurate pricing, prevention and portfolio steering.

Embedded and API-native distribution models

API-first platforms allow insurance to be embedded seamlessly into mobility, retail, banking and partner journeys, shifting growth from standalone products to contextual, usage-based and event-driven protection.

Emotion- and behavior-aware customer engagement

Affective analytics and behavioral insights personalize communication, prevention nudges and claims interactions, strengthening trust, engagement and long-term retention in increasingly digital journeys.

Trusted AI, cybersecurity and decision governance

Explainable models, AI lifecycle governance, deepfake detection and adaptive cyberdefenses ensure underwriting and claims decisions remain auditable, fair and resilient as automation scales.

Informed, modular and ecosystem-ready infrastructure

Hybrid cloud-edge architectures, interoperable data layers and modular services allow insurers to scale AI, integrate partners rapidly and respond dynamically to risk events without sacrificing reliability.

Sovereign computing, data residency and post-quantum readiness

In-region processing, confidential computing, diversified infrastructure providers and early post-quantum preparation strengthen compliance, resilience and operational continuity in a fragmented regulatory landscape.

Industry use case

Geospatial risk twin and trusted AI

A global insurer creates a continuously updated geospatial risk twin that integrates satellite imagery, 3D property models, hazard maps, historical loss data and exposure insights. Before catastrophic events, teams simulate concentration risk, loss severity and reinsurance strategies. During and after large-scale loss events, trusted AI models refresh the twin using imagery, IoT and claims data, proposing best-next actions from triage adjustments to parametric payouts.

Results

- Faster and more accurate loss estimation
- Earlier customer support
- Reduced fraud
- More resilient portfolio decisions
- Transparent reporting for regulators and reinsurers

Strategic priorities

01

Build risk intelligence systems for resilient portfolios

Integrate geospatial, IoT, telematics, biosignal and digital-twin data to guide underwriting, pricing, accumulation and capital allocation.

02

Industrialize responsible agentic AI across the value chain

Deploy governed XAI frameworks for underwriting, claims, fraud and customer operations.

03

Shift to emotion-driven, prevention-first and embedded models

Use behavioral analytics, biometrics, telematics and embedded APIs to deliver personalized prevention and seamless in-journey protection.

Automotive

Automotive enters 2026 under intense pressure to electrify profitably, scale software-defined vehicles (SDVs), secure battery and semiconductor supplies, and navigate geopolitical fragmentation across manufacturing and technology stacks.

Electric vehicle (EV) adoption continues to rise, yet margins lag due to battery costs, shifting subsidies, charging-network immaturity and strong global competition, particularly from Chinese original equipment manufacturers (OEMs).

Vehicles are becoming always-connected digital platforms, creating new expectations for cybersecurity, data sovereignty and software quality. New mobility behaviors — like subscriptions, shared use and home-charging patterns — further transform revenue models.

The next wave of differentiation will come from autonomy, intelligent factories, informed energy ecosystems, sustainable materials and circular EV platforms, requiring synchronized roadmaps across hardware, software, energy and ecosystem partnerships.



Main industry challenges



EV profitability, battery economics and circularity

OEMs must improve unit economics through platform reuse, diverse chemistries, recycling, second-life flows and transparent total cost of ownership (TCO) models.



SDV reliability and over-the-air (OTA) quality

Software-rich vehicles require robust DevOps/MLOps pipelines, safety cases, silicon attestation and update quality to preserve trust and revenue.



Semiconductor, supply chain and geopolitical resilience

Diversified suppliers, codesigned chips, localized manufacturing and risk-mitigated supply networks reduce exposure to critical chokepoints.



Charging infrastructure, energy integration and CX

Reliable charging, tariff intelligence and vehicle-to-everything (V2X)-enabled energy services are essential for customer satisfaction and EV adoption.



Cybersecurity, safety and data governance

United Nations Economic Commission for Europe, International Organization or Standardization/Society of Automotive Engineers (UNECE) regulations, ISO/SAE standards, data-consent rules and secure-by-design architectures drive continuous monitoring and compliance.

Trend relevance on the automotive value chain

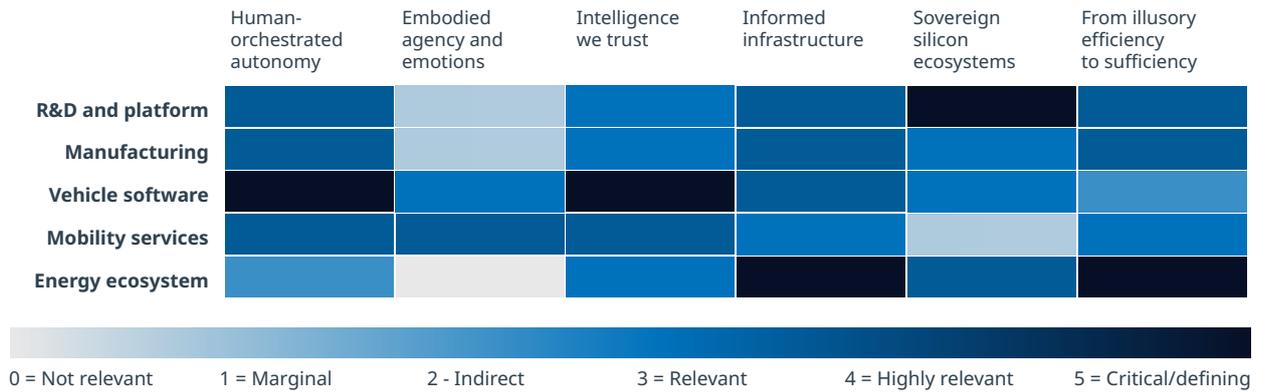


Figure 16: Automotive value chain: Trend relevance heatmap

Transformation drivers

Agentic vehicle, factory and fleet operations

AI agents coordinate production planning, quality control, OTA rollouts and fleet-behavior tuning, enabling faster response and continuous optimization while keeping human oversight for safety and compliance.

High-fidelity digital twins across product and lifecycle

Vehicle, plant, battery and fleet digital twins simulate performance, defects and energy usage, reducing launch risk, accelerating development and improving lifecycle cost and reliability.

Software-defined vehicle platforms and continuous evolution

Industrialized DevOps/MLOps, modular SDV architectures and secure OTA pipelines enable frequent feature updates and monetization without compromising safety, quality or customer trust.

Emotion- and context-aware in-vehicle experiences

Driver-state sensing, affective AI and contextual assistants personalize interactions, reduce distractions and increase the adoption of digital services and subscriptions.

Trusted software, cybersecurity and functional safety

Secure-by-design vehicle stacks, traceable software lifecycles and safety-critical validation protect vehicles, customers and brands as connectivity and autonomy increase.

Informed energy and charging ecosystems

Real-time analytics coordinate vehicles, charging infrastructure, grid signals and home energy systems to improve availability, reduce cost and enhance the ownership experience.

Sovereign and resilient semiconductor and computing stack

Codesigned automotive-grade silicon, dual sourcing and localized computing reduce exposure to supply shocks and geopolitical constraints while improving performance-per-watt for SDVs.

Circularity and sustainable material systems

Product passports, recycling intelligence and AI-driven reuse pipelines support decarbonization, material stability and ESG credibility.

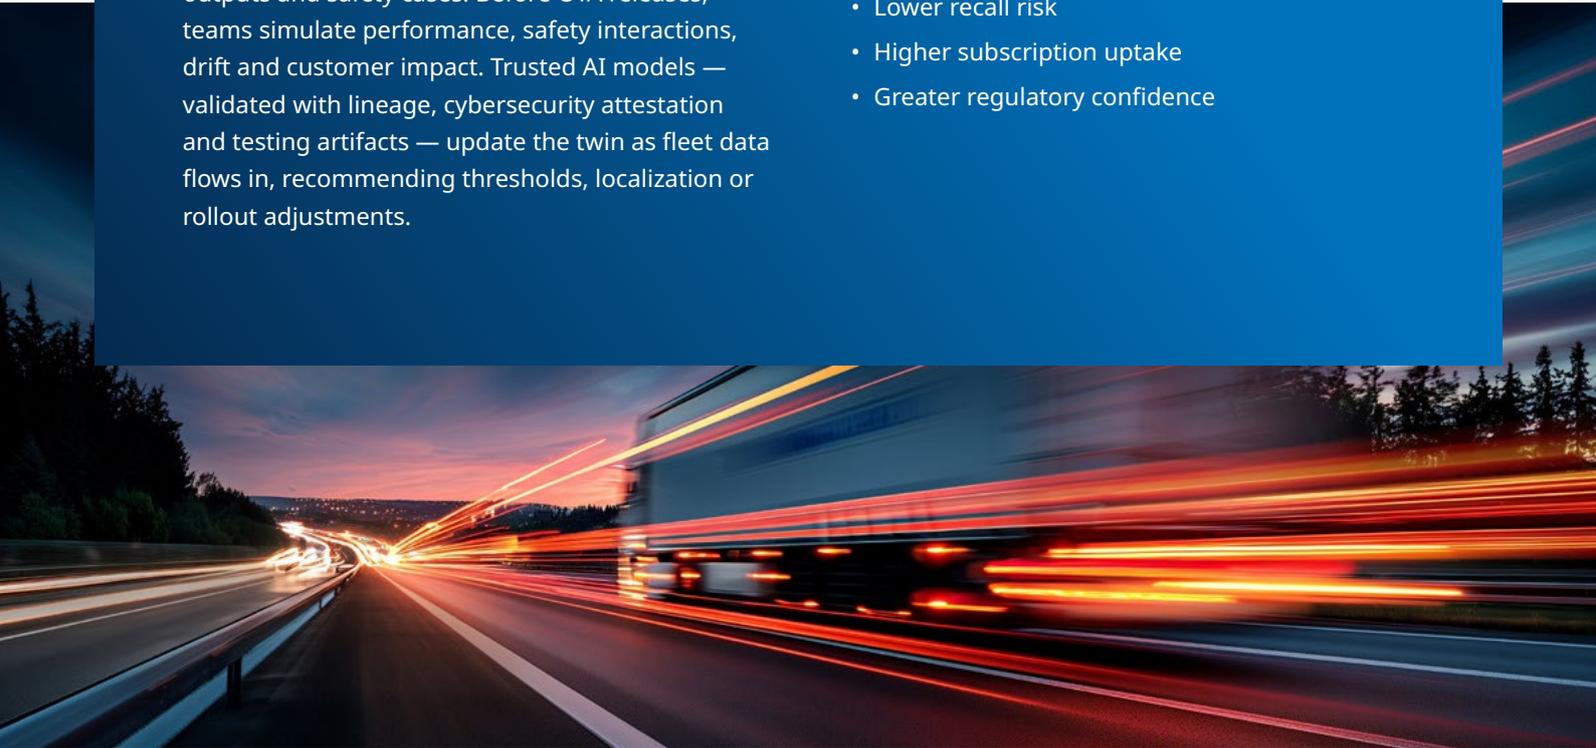
Industry use case

Software twin and trusted OTA

A global OEM builds a software twin that integrates SDV architectures, field data, sensor logs, simulation outputs and safety cases. Before OTA releases, teams simulate performance, safety interactions, drift and customer impact. Trusted AI models — validated with lineage, cybersecurity attestation and testing artifacts — update the twin as fleet data flows in, recommending thresholds, localization or rollout adjustments.

Results

- Faster release cycles
- Fewer defects
- Lower recall risk
- Higher subscription uptake
- Greater regulatory confidence



Strategic priorities

01

Industrialize SDV and agentic operations across product and production

Unify engineering, manufacturing, quality and OTA workflows with digital twins and governed autonomy.

02

Build a resilient, sovereign stack of energy-computing materials

Secure batteries, semiconductor resilience, sovereign-cloud data flows and integrated energy ecosystems for stability and performance.

03

Shift to lifecycle-centric, experience-driven revenue models

Leverage telemetry, emotional intelligence, smart charging and transparent subscriptions to maximize lifetime value.



Telecommunications

Telecommunications enters 2026 under simultaneous growth, cost and sovereignty pressures.

While traffic continues to surge, revenues remain flat, creating a structural race to lower cost-per-bit. Operators must balance 5G standalone deployment, fiber expansion, open radio access network (RAN) maturity and telco-cloud migrations amid sustained capex.

At the same time, enterprises demand private 5G, multiaccess edge computing (MEC), network slicing, IoT integration and API-based assurance — requiring operators to evolve from connectivity providers to outcome-based service platforms.

Energy price volatility, sustainability expectations and geopolitical pressures on spectrum, cloud partnerships and silicon supply further shift telcos to operating as critical-infrastructure providers. Customers expect frictionless, app-first interactions, while developers anticipate programmable, open networks.

These forces accelerate the transition from telco to techco, with software, AI, automation and platform ecosystems becoming the core of competitiveness.

Main industry challenges



Monetizing 5G, MEC and APIs

Operators must turn network slicing, private wireless networks and programmable APIs into scalable enterprise and developer offerings.



Network economics, automation and unit-cost reduction

AI for IT operations (AIOps), autonomous networks and cloud-native cores are essential to reversing cost-per-bit erosion while improving reliability.



Customer experience, embedded journeys and fair bundles

Differentiation increasingly depends on transparent plans, sentiment-aware interactions and ecosystem bundles — including over-the-top media (OTT) content and selective traffic subsidization — that strengthen loyalty and reduce churn.



Security, resilience and data sovereignty

Expanding attack surfaces and new regulations demand zero trust, in-region processing, supply chain assurance and operational continuity.



Semiconductor, cloud and vendor dependence

Growing influence from hyperscalers and chipmakers requires silicon diversity, sovereign-cloud options and Open RAN maturity.

Trend relevance on the telco value chain

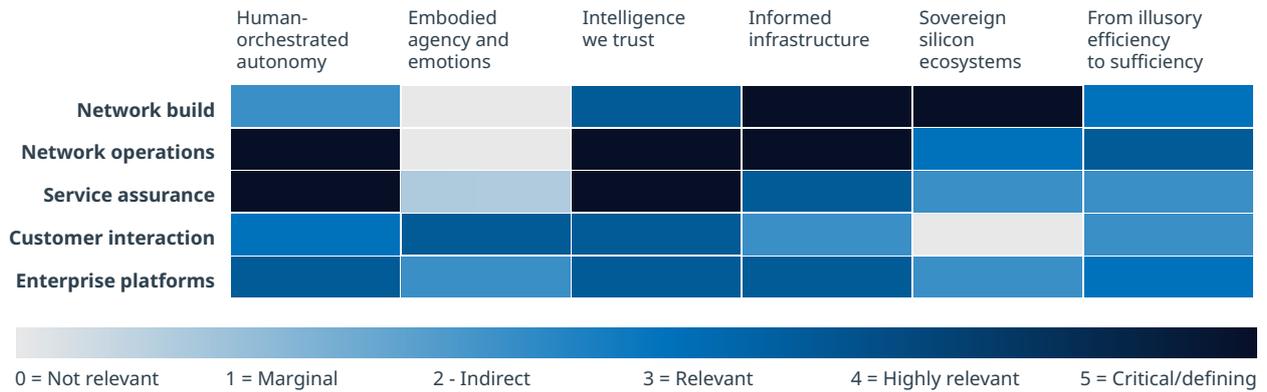


Figure 17: Telco value chain: Trend relevance heatmap

Transformation drivers

Agentic network operations and human-guided automation

AI-driven, closed-loop control automates assurance, optimization and recovery across RAN, transport and core networks, improving reliability and speed while keeping humans accountable for safety, policy and escalation.

Programmable networks, APIs and outcome-based services

Network slicing, private wireless, MEC and standardized APIs expose network capabilities as programmable services, enabling enterprises and developers to consume assured latency, reliability and security on demand.

Informed, real-time network infrastructure

Continuous telemetry, digital twins, AIOps and edge-cloud intelligence turn networks into self-monitoring systems that optimize performance, energy use and rollout efficiency in real time.

Customer experience, embedded journeys and ecosystem bundles

App-first care, sentiment-aware interactions and ecosystem bundles — including OTT content and service integration — strengthen customer loyalty and reduce churn in consumer and enterprise segments.

Security, resilience and sovereign network operations

Zero trust architectures, in-region processing, supply chain assurance and tested continuity plans protect expanding IT/OT attack surfaces and position telcos as operators of critical national infrastructure.

Cloud, silicon and vendor diversification

Open RAN, multicloud strategies and diversified silicon sourcing reduce dependency on hyperscalers and chipmakers while preserving cost control, performance and long-term strategic autonomy.

Industry use case

Closed-loop AI assurance for enterprise SLAs

A European operator deploys AI-driven, closed-loop assurance for private 5G and MEC customers. Real-time telemetry from RAN, transport, and edge is fused with SLA parameters to detect anomalies, predict congestion, adjust network slice configurations and reroute traffic autonomously. When thresholds are at risk, the system notifies customers, triggers corrective actions and logs steps for audit.

Results

- Higher SLA attainment
- Lower mean time to repair (MTTR)
- Reduced churn
- Improved energy efficiency
- Full transparency for regulators and enterprise clients

Strategic priorities

01

Industrialize autonomous networks with trusted AI

Deploy AIOps, digital twins and closed-loop assurance across network domains with explainability, lineage and safety frameworks.

02

Shift from connectivity to outcome-based services and APIs

Build automated productization and delivery platforms for network slicing, MEC, private 5G and Open Gateway APIs, exposing programmable network capabilities with quality-of-service guarantees.

03

Run networks as sustainable, sovereign infrastructure

Lower energy per bit, regionalize critical workloads, diversify silicon, and ensure sovereign cloud-edge alignment for regulatory compliance.



Manufacturing

Manufacturing enters 2026 in a phase of structural volatility marked by supply chain shocks, geopolitical fragmentation, inflationary pressure, labor scarcity and rising energy costs.

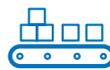
Digital transformation is essential but often slowed by heterogeneous systems, legacy data and uneven scaling across multiplant networks. At the same time, rapid advances in physical AI, agentic intelligence, real-time decision-making and intelligent automation are reshaping factory operations — moving them toward adaptive, resilient and more human-centric production models.

Main industry challenges



Scaling smart factory ROI across heterogeneous plants

Achieving productivity, quality, energy reduction and autonomy across diverse sites requires standardized data, toolchains and operating models.



Supply chain resilience under volatility

Persistent disruptions and supplier concentration demand multisourcing, real-time visibility and agile responses that balance cost and working capital.



Workforce transformation and human-machine collaboration

Skills gaps in automation and data require reskilling and new human-machine collaboration models as autonomy expands.



Data foundation and cyber-physical trust

Scaling AI depends on unified, governed data and robust information technology/operational technology (IT/OT) security for traceability, uptime and compliant model deployment.



Sustainability, energy and capital efficiency

Sufficiency, circularity and decarbonization targets influence product, resource and investment decisions as cost and carbon performance converge.

Trend relevance on the manufacturing value chain

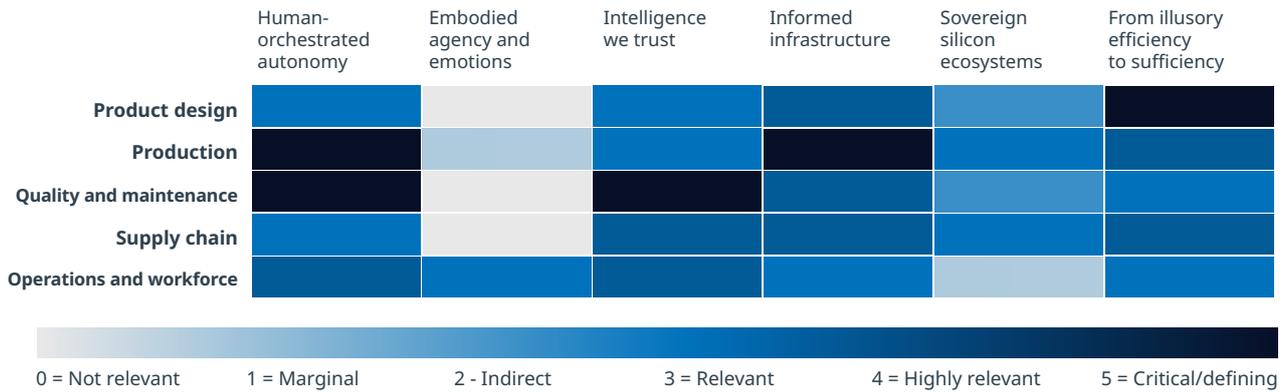


Figure 18: Manufacturing value chain: Trend relevance heatmap

Transformation drivers

Agentic factory operations and human-guided automation

AI agents orchestrate planning, production, quality and response across machines, lines and plants, enabling faster adaptation to disruptions while keeping humans accountable for safety, ethics and performance trade-offs.

Unified industrial data foundation and interoperability

Vendor-neutral, real-time data layers connect machines, OT systems, manufacturing execution systems and manufacturing operations management (MES/MOM) platforms and applications, creating a scalable foundation for analytics, digital twins and cross-plant replication of value.

Physical AI, robotics and flexible automation

AI embedded in machines, robots and autonomous cells enables adaptive production layouts, small-batch responsiveness and safer human-machine collaboration without constant line reconfiguration.

Informed, real-time manufacturing infrastructure

Edge-cloud architectures, continuous telemetry and operational digital twins optimize throughput, energy use, maintenance and launch stability across plants and supply networks.

Trusted AI, cyber-physical security and governance

Explainable models, IT/OT security-by-design, traceability and policy-driven access controls protect production integrity, quality and uptime as autonomy and connectivity scale.

Sufficiency-driven production and circular operations

Rightsized capacity, modular platforms, material reuse and energy-aware optimization shift manufacturing from pure efficiency to resilient, sustainable and capital-efficient operations.

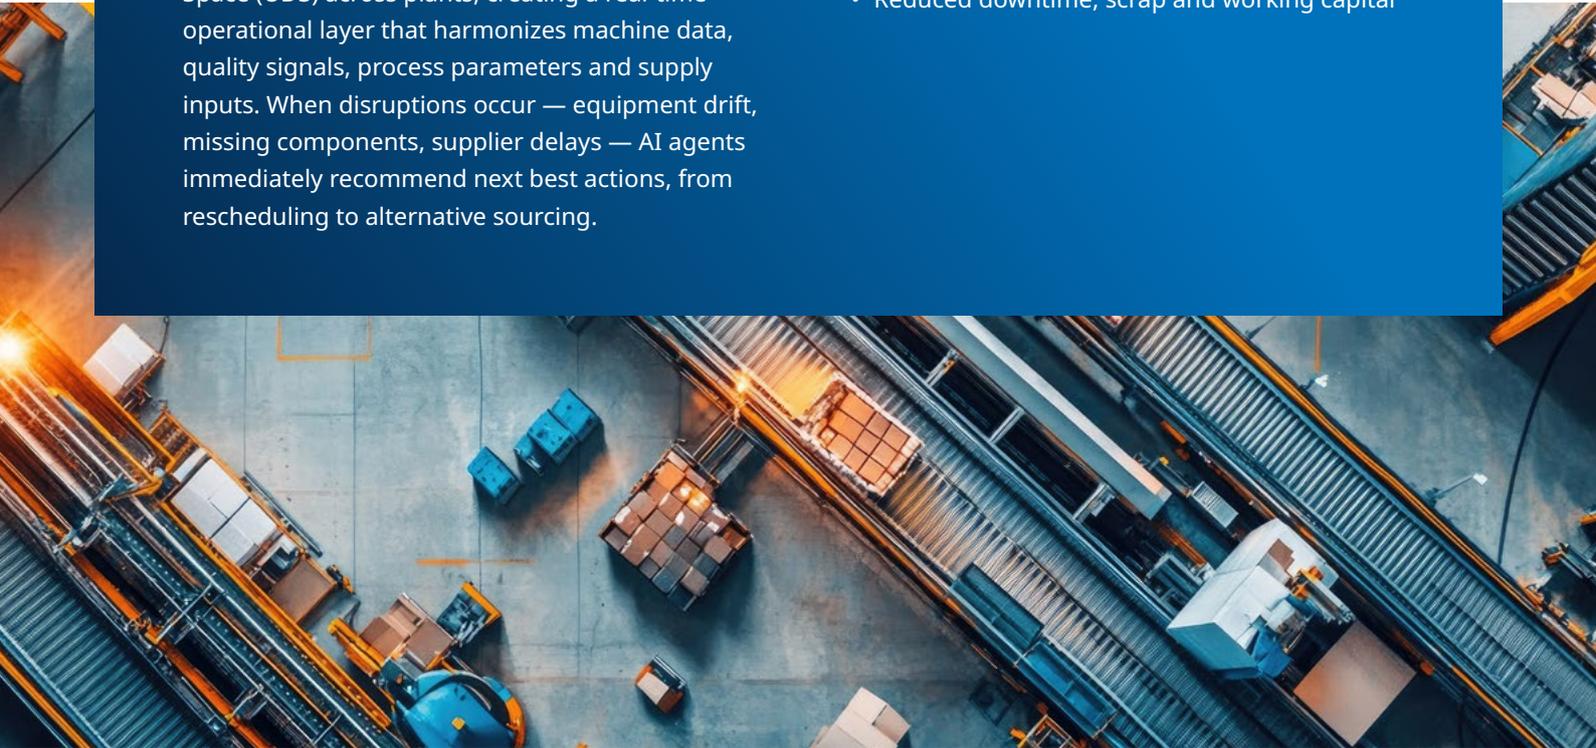
Industry use case

Real-time disruption response with AI and Unified Data Space

A global CPG manufacturer deploys a Unified Data Space (UDS) across plants, creating a real-time operational layer that harmonizes machine data, quality signals, process parameters and supply inputs. When disruptions occur — equipment drift, missing components, supplier delays — AI agents immediately recommend next best actions, from rescheduling to alternative sourcing.

Results

- Disruption handling shifts from manual firefighting to continuous, automated resilience
- Reduced downtime, scrap and working capital



Strategic priorities

01

Build a scalable data foundation (Unified Data Space with governance)

Implement a unified, vendor-neutral data layer to unlock cross-plant replicability and enable predictive maintenance, digital twins and agentic control.

02

Operationalize agentic and autonomous decision-making

Deploy AI agents for planning, scheduling, maintenance and supply chain response while ensuring human governance and safety controls.

03

Design for flexible production with hybrid robotics

Combine human operators, cobots and humanoids to increase agility, support labor shortfalls and reduce downtime.



Banking

Banking enters 2026 under structural pressure on cost, trust, compliance and growth.

Competitive dynamics intensify as fintechs scale, embedded finance expands, digital wallets proliferate and payments shift to instant and account-to-account rails.

Regulators introduce stringent frameworks — the Digital Operations Resilience Act (DORA), Network and Information Systems Directive 2 (NIS2), General Data Protection Regulation (GDPR), Framework for Financial Data Access (FIDA) in the EU, and anti-money laundering (AML) packages, among others — requiring banks to operate cloud, data, AI and third-party ecosystems as regulated critical infrastructure.

Customer expectations for instant transactions, transparent fees, personalized advice and emotionally intelligent servicing continue to rise. Meanwhile, climate disclosures, taxonomy alignment, liquidity optimization and deposit stability create additional macro uncertainty.

As central bank digital currencies (CBDCs) and early-stage quantum and photonic infrastructures emerge, banks must modernize core platforms, trusted AI, payments, risk and ecosystems while demonstrating operational resilience and sovereignty.

Main industry challenges



Core, payments and digital infrastructure modernization

Banks must modernize core systems, payment rails, cloud architectures and data foundations to support real-time services and the digitalization of money, including digital currencies, wallets, identities and ledger-based settlement.



Cybersecurity, fraud and financial crime complexity

Escalating threats across identity, payments and APIs require continuous authentication, network analytics and explainable, auditable models.



Regulatory scrutiny, sovereignty and operational resilience

Regulators demand audited controls, in-region processing, transparent lineage and rigorous third-party governance.



Trusted AI, personalization and ethical decision-making

Personalized advice, underwriting and servicing rely on governed AI with strict consent, fairness, suitability and bias-mitigation controls.



Go-to-market transformation and margin management

Open finance, embedded channels and changing customer behavior require new pricing engines, product architectures and customer lifetime value (CLV)-centric strategies.

Trend relevance on the banking value chain

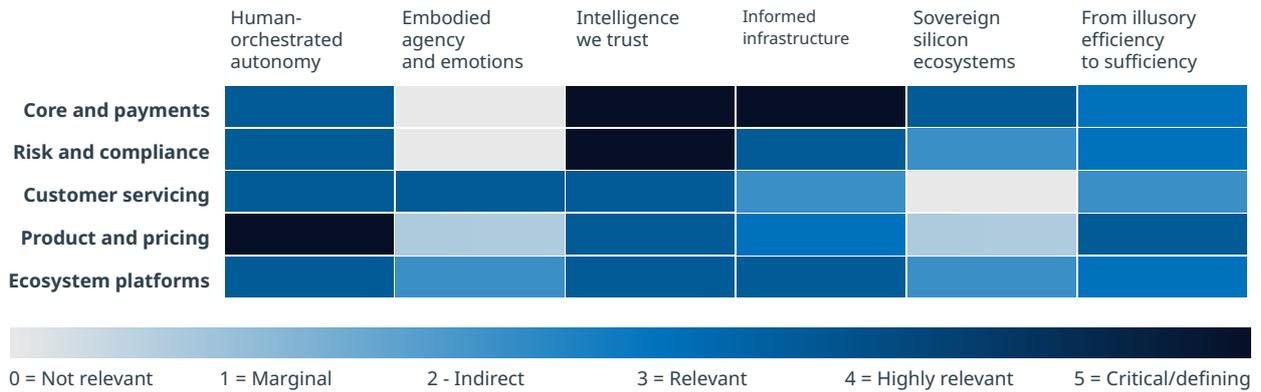


Figure 19: Banking value chain: Trend relevance heatmap

Transformation drivers

Agentic banking operations and human-guided automation

AI agents orchestrate multistep processes across payments, onboarding, servicing, finance, and operations, increasing speed and efficiency while preserving human approval, segregation of duties and auditability.

Real-time core, payments and ledger modernization

Event-driven cores, instant payment rails and modern ledger capabilities enable real-time processing, programmable money use cases and scalable support for digital currencies, wallets and settlement innovation.

Trusted AI for decision-making, risk and personalization

Governed AI models power underwriting, fraud prevention, collections and next best actions with explainability, fairness controls, lineage and continuous monitoring to meet regulatory and customer trust expectations.

Digital identity, authentication and consent fabric

Continuous authentication, strong digital identity and consent management underpin secure access to accounts, APIs, wallets and embedded journeys across channels and partner ecosystems.

Informed, resilient and sovereign banking infrastructure

Streaming architectures, observability-first operations, multicloud failover and in-region processing ensure always-on services, regulatory compliance and rapid recovery for critical banking workloads.

Embedded and platform-based distribution models

API-first architectures allow banks to distribute payments, credit and services through third-party platforms and ecosystems, shifting growth to contextual, embedded and usage-driven journeys.

Cybersecurity and financial crime defense at scale

Advanced fraud analytics, graph-based anti-money laundering (AML), behavioral signals and coordinated cyberdefenses protect high-speed payment flows and open interfaces while keeping investigations auditable and regulator-ready.

Industry use case

Agentic fraud and financial crime control for instant payments

A Tier 1 bank deploys a real-time fraud and AML control fabric that consumes streaming payments, device fingerprints, behavioral patterns and customer context. AI agents classify risk, orchestrate next best actions — from step-up authentication to sanctions checks — and auto-generate cases with entity linkages and recommended actions.

Results

- Lower fraud losses and false positives
- Improved customer experience
- Reduced investigation costs
- Ability to meet stringent regulatory expectations

Strategic priorities

01

Industrialize agentic and trusted AI across core, risk and operations

Deploy governed AI agents with strict scopes, lineage, approvals, and privacy and fairness controls across servicing, payments, AML, know your customer (KYC) and underwriting.

02

Run banking on informed, sovereign, always-on infrastructure

Adopt streaming-first architectures, multicloud resilience and sovereign computing zones to meet operational resilience and regulatory compliance mandates.

03

Shape the new financial rails (instant, open, CBDC, tokenized assets)

Modernize payments with ISO 20022 intelligence, real-time fraud controls, open-finance orchestration and readiness for CBDCs and tokenized assets.

Pharma

Pharma enters 2026 under converging pressures: a major patent cliff, tougher pricing and access regimes (the Inflation Reduction Act (IRA) in the US and the Health Technology Assessment (HTA)), and increasing complexity in biologics, cell and gene therapy (CGT) and antibody-drug conjugate (ADC) modalities.

Expectations for digital, data and ESG performance continue to rise. At the same time, pipelines shift to high-science, high-cost therapies, while payers demand real-world value and regulators intensify their scrutiny of AI use, data integrity and GxP (Good [x] Practices) compliance.

Operations must become more agile and resilient as multimodal manufacturing, fragile supply chains and stringent quality requirements coincide with rising cost and access pressures.



Main industry challenges



Pipeline productivity and evidence-driven decision-making

High-science modalities and rising development costs increase the need for earlier predictive insights, faster go/no-go decisions and targeted investments.



Modernizing clinical trials and patient access

Decentralized and digital trial models, diversity requirements and new regulatory expectations affect recruitment, design and real-world evidence generation.



Manufacturing agility and modality scale-up

Biologics, CGT, ADCs and multimodal networks require flexible, validated production supported by digital twins, quality intelligence and technology transfer excellence.



Trusted AI, data governance and regulatory compliance

Scaling AI across R&D, clinical and manufacturing demands rigorous model lineage, consent, privacy, explainability and secure multiregion data flows.



Supply chain, sustainability and global fragmentation

Active ingredient shortages, geopolitical pressures, decarbonization mandates and ESG disclosures require transparent, resilient global operations.

Trend relevance on the pharma value chain

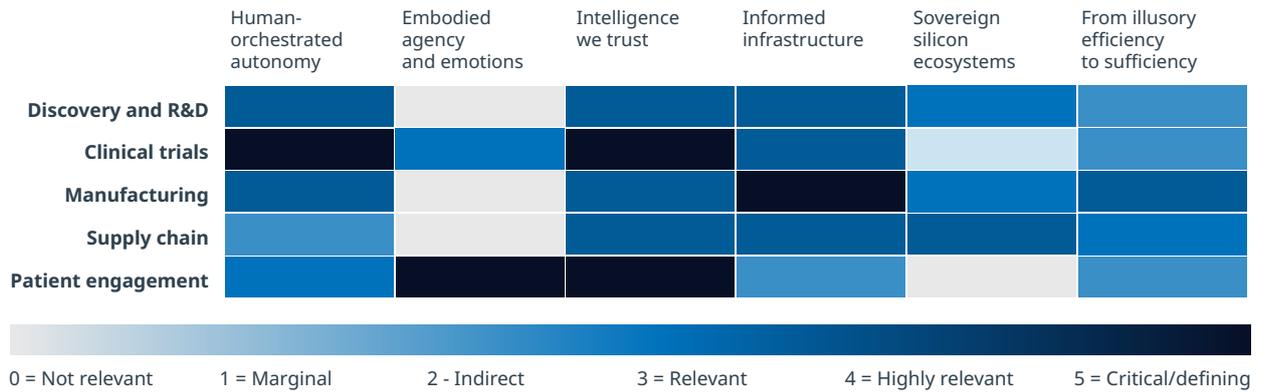


Figure 20: Pharma value chain: Trend relevance heatmap

Transformation drivers

Agentic R&D, clinical and operations orchestration

AI agents coordinate discovery, trial operations, safety and manufacturing workflows across multiple steps, accelerating decisions while preserving human oversight, quality accountability and regulatory control.

Clinical-trial digital twins and real-world evidence integration

Virtual trial environments unify protocol designs, historical trial data, real-world evidence (RWE), and site performance to simulate outcomes, anticipate risk and optimize enrollment, diversity and execution in real time.

Adaptive, modality-aware manufacturing systems

Digital twins, advanced analytics and flexible automation enable rapid scale-up and transfer across biologics, CGT, ADCs and multimodal networks while protecting GMP compliance and yield.

Trusted AI, data governance and regulatory-grade models

Explainable models, lineage tracking, validation artifacts, and lifecycle governance ensure AI can be safely scaled across R&D, clinical, pharmacovigilance and manufacturing under GxP requirements.

Informed, resilient and sovereign pharma infrastructure

Hybrid cloud-edge architectures, real-time observability and in-region processing support secure data sharing, resilient operations and compliant collaboration across global R&D and supply ecosystems.

Patient-centric digital engagement and evidence generation

Digital approaches such as eConsent, eSource, wearables and digital biomarkers enable continuous patient interaction and evidence capture, improving access, adherence and insight in both the development and post-market phases.

Industry use case

Clinical-trial twin and trusted AI

A global pharma company builds a clinical-trial digital twin, integrating protocol design, historical trial data, real-world evidence and site performance. Before first-patient-in, teams simulate enrollment curves, diversity targets, site risk and protocol deviations. Trusted AI models — validated with full lineage — continuously update the twin as real-world data flows in, recommending actions such as activating backup sites or adjusting visit schedules.

Results

- Faster database lock
- Improved diversity
- Fewer amendments
- Better inspection readiness
- Full transparency for regulators and partners

Strategic priorities

01

Industrialize trusted AI across the value chain

Establish a unified model governance and model risk management (MRM) framework covering discovery, trials, safety, manufacturing and forecasting.

02

Build validated, informed-infrastructure and data products

Invest in digital twins, cross-IT/OT observability, and certified data products to accelerate evidence generation and decision quality.

03

Redesign operating models around access, sustainability and talent

Align market access, medical and real-world-evidence functions to value-based agreements; embed ESG into supplier governance; and build digital and AI capabilities for chemistry, manufacturing and controls (CMC), pharmacovigilance (PV) and clinical operations.

From insight to action

The convergence map

The convergence map illustrates how the six macro trends intersect to form a unified transformation ecosystem — the blueprint for the future enterprise. It reveals dynamic connections where technological innovation, digital sovereignty and sustainability both overlap and reinforce each other, becoming inseparable drivers of progress.

By viewing these six macro trends as interdependent components of mass intelligence, organizations can move from isolated initiatives to coherent strategies that align technology, purpose and long-term resilience.

Convergence map of the six macro trends

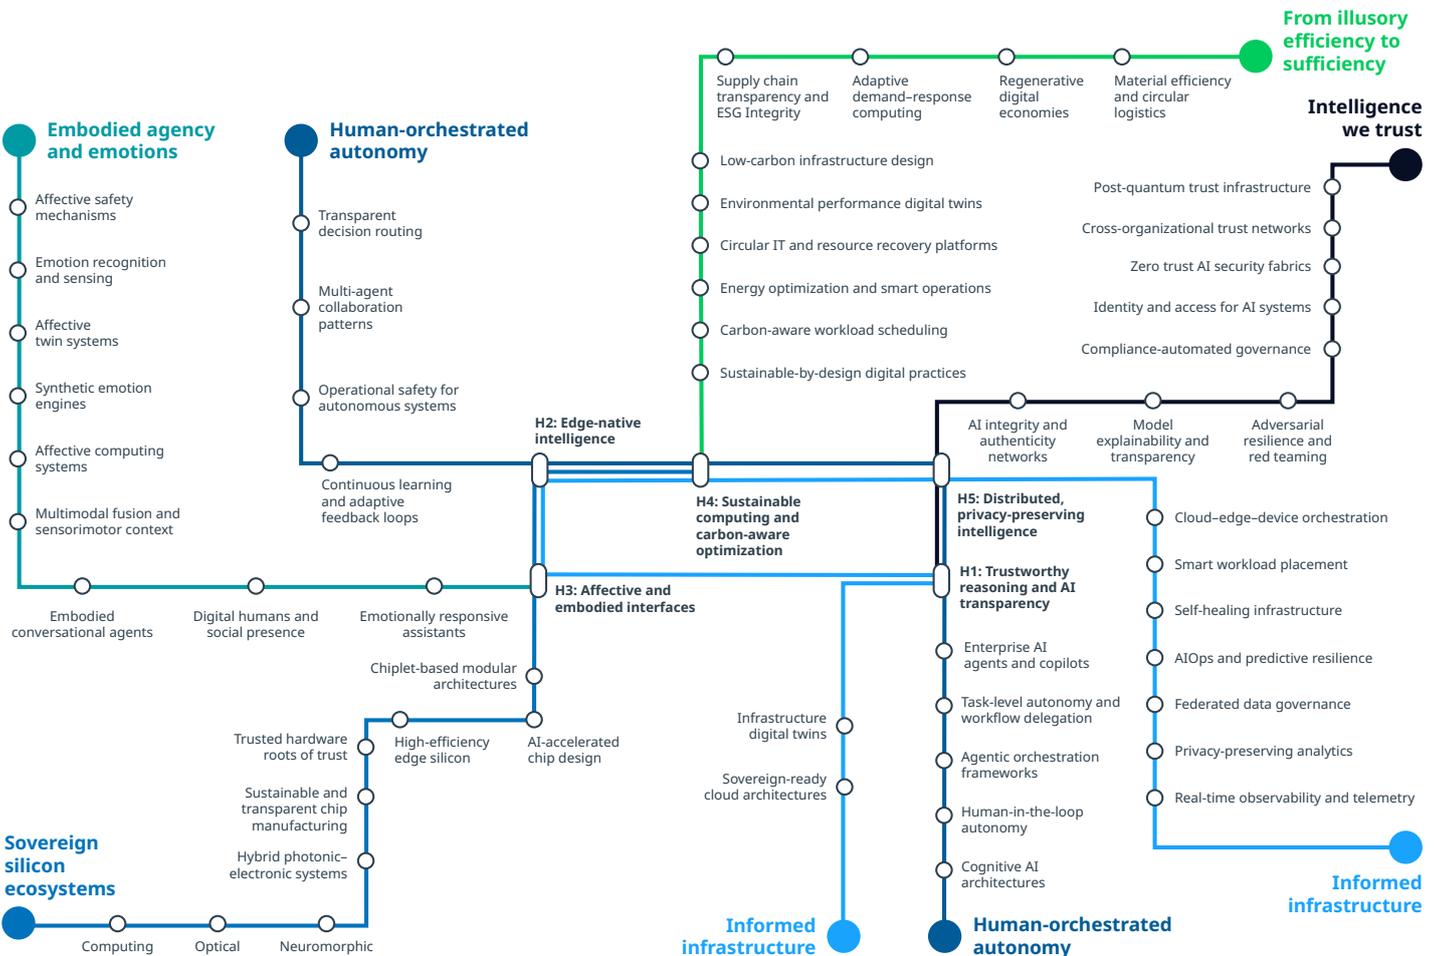


Figure 21: Convergence map of the six macro trends

Strategic capabilities

To stay competitive in this increasingly complex landscape, organizations must develop strategic capabilities that both elevate human potential and strengthen technological foundations.

This begins with augmenting human intelligence, shifting from automation to human-AI symbiosis. At the same time, they must design for ethical growth, making transparency, trust and regenerative principles the default. Building connected foundations is equally essential, ensuring interoperable, fluid and sovereign infrastructures. Organizations also need to build foresight-ready cultures that embed anticipation, experimentation, capability building and agility in all activities. Finally, they must accelerate collaboration ecosystems, forging partnerships beyond traditional boundaries to create adaptive value networks.

An informed approach to navigating potential

While the possibilities and potential of the technology macrorends we've explored in this report are exciting, we must remain alert to their potential risks.

Uncertainties will always persist, but striving to comprehend future scenarios can decrease the likelihood of facing unexpected challenges and blind spots.

As these six macrorends converge, they will redefine the way we work, shifting technology from automating tasks to amplifying human capability, and enabling people and intelligent systems to co-create value in more adaptive, ethical and purpose-driven ways.

NTT DATA stands ready to help

Our services include business and technology consulting, data and artificial intelligence, industry solutions, as well as the development, implementation and management of applications, infrastructure and connectivity.

We are also one of the leading providers of digital and AI infrastructure in the world.

NTT DATA is part of NTT Group, which invests over \$3 billion each year in R&D to help organizations and society move confidently and sustainably into the digital future.



About the research

The methodology behind NTT DATA's Technology Foresight 2026 combines rigorous research practices with unique internal resources to explore and refine emerging macrorends in technology.

Drawing from NTT Group's extensive investments in basic and applied research, our team benefits from direct access to cutting-edge developments in a wide array of fields, enabling a deep understanding of technological advancements and their potential applications.

A key component of the process is the use of an internal market intelligence platform which aggregates and analyzes data from diverse sources, including market forecasts, analyst insights and public information on investments and partnerships. This automated, data-driven approach ensures that trend selection is grounded in objective evidence.

The methodology is further enhanced by consultation with a global network of technology and business experts in NTT Group. These specialists contribute localized and domain-specific insights, refining macrorends through a collaborative, multidisciplinary lens.

Finally, techniques from strategic foresight are employed to map these macrorends and use cases to future scenarios, examining their implications for individuals, businesses and society. This structured, iterative process ensures that the identified macrorends are both technologically robust and contextually relevant.

Contributors

Editorial:

Hisashi Yoshida

Cynthia Yumiko Staub

Rena Oi

Valentina Contini

Oliver Köth

Alberto Acuto

Contributors:

Raoul Heese

Iris Pfeifer

Amulya Bhatia

Diana Hauser

Bettina Wegner

Doris Bruder

Volker Ganz

Markus Lunz

Johannes Wetterich

Dennis Tischer

Afkar Alam Khan

Tobias Brecht

Christian Koch

Filippo Capocasale

César Zayas

Gadhu Sundaram

Mireia Vilaplana

Victor León Marambio

Pablo Sáez Nunez

Enno Kätelhön

Cornelius Walter

Bernd Pape

Andreas Schlüter

Tobias Schmidt

Martin Pölzl

Andreas Möbs

Rouven Steibli

Ruzica Lovric

Aylin Lindmayer

Carlos Ordonez Sánchez

Richard Calvo Lopez

Rahul Maheshwari

Amit Unde

Maria Cione

Carlos Galve Pellicero

Natasha Horwitz

Kate Owendale

Retha Duvenage

Jessica Brink

Anneli Brooks-Smith

Stefania Vita

Sara Alvarellos Navarro

Jose Carlos Chavez Peinado

Raquel Serradell Muñoz

Aranxa Luisa Pino Baez

Lucia Ciordia Navarro

Paula van Huyssteen

Jessica Acar

Georg Kerpert

Daniela Florio

Visit nttdata.com to learn more.

NTT DATA is a \$30+ billion business and technology services leader in AI and digital infrastructure. We accelerate client success and positively impact society through responsible innovation. As a Global Top Employer, we have experts in more than 70 countries. NTT DATA is part of NTT Group.



