Looking Glass

Bringing tech-led business changes into focus

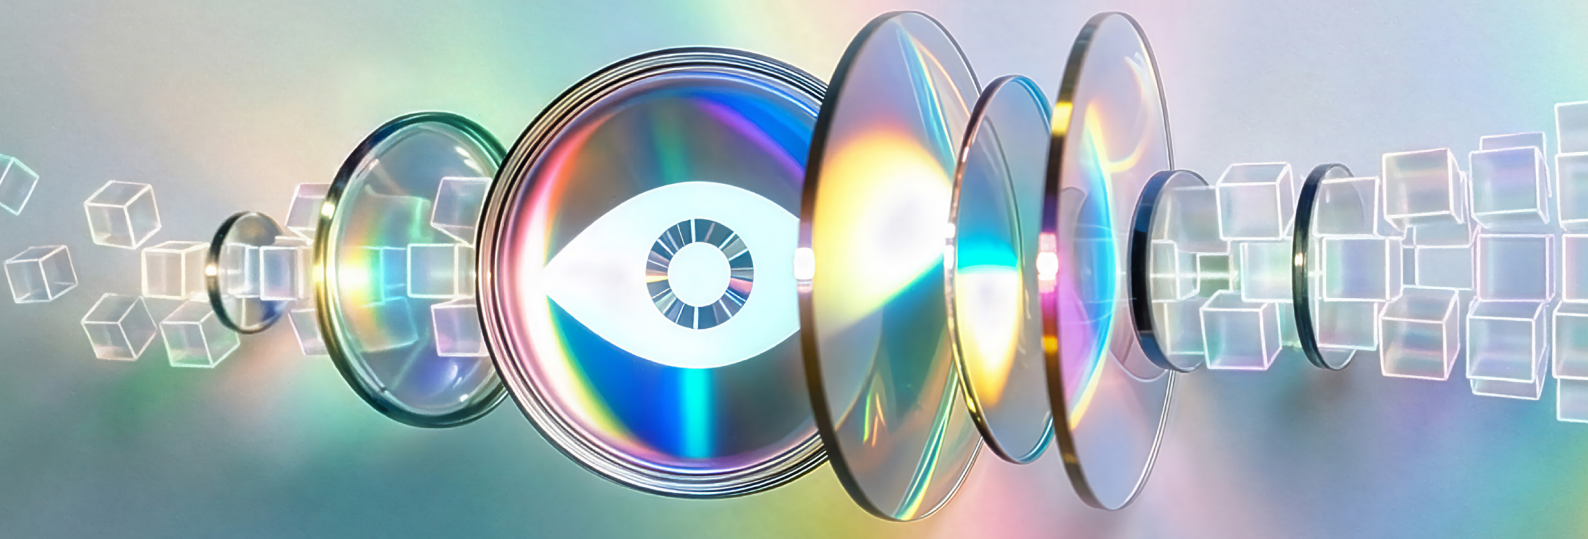

thoughtworks

Design. Engineering. AI.

2026

Introduction

Each year, Thoughtworks' Looking Glass report looks beyond individual technologies to examine the forces reshaping how organizations build, run and evolve their technology estates. The aim is not to predict the future, but to help leaders make sense of the changes already underway — and to understand which of them will matter most in practice.

This year's trends reflect a moment of transition. Long-running shifts around platforms, data, security and experience design are converging with rapid advances in AI. The result is not a single disruptive technology, but a reconfiguration of how technology creates value across the enterprise. Systems are becoming more adaptive. Interactions more intent-driven. Governance more embedded in day-to-day delivery.

Across the Looking Glass, we explore what this means in concrete terms: how enterprises are rebuilding their core foundations, rewiring workflows to support greater autonomy and reimagining the role technology plays in customer experience, decision-making and operations. AI features prominently, but always in context — as one of several forces accelerating change, rather than a solution in isolation.

As ever, this report is grounded in what we see working with clients around the world. It is intended as a practical guide to navigating the year ahead, with a focus on technologies that deliver real, durable value.

—

Rachel Laycock
Chief Technology Officer, Thoughtworks

1

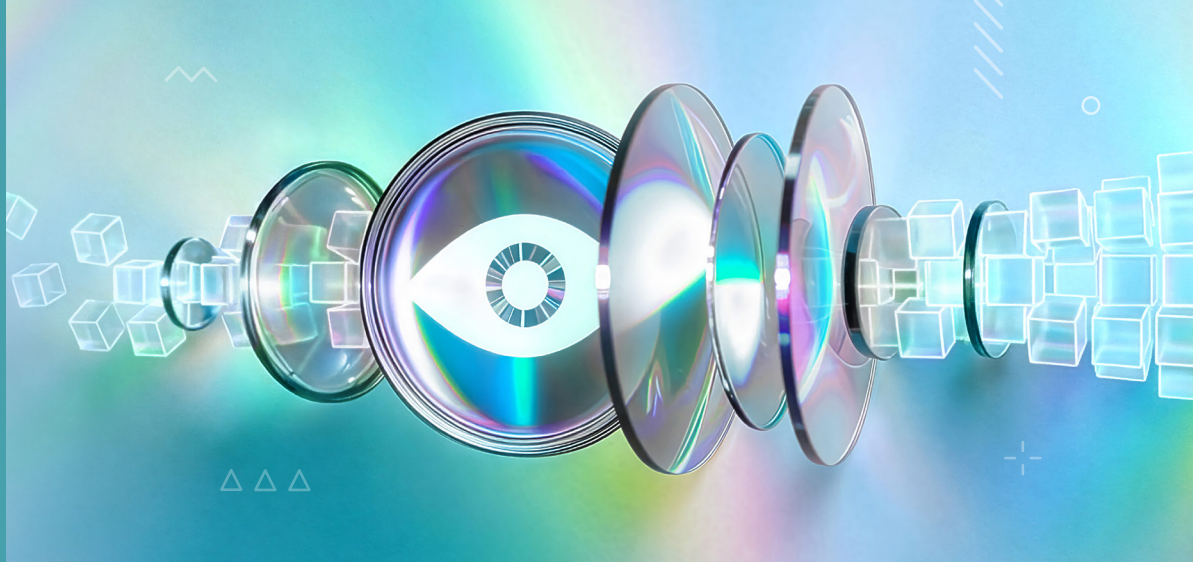

AI and software delivery

Time to rebuild core systems

Adoption of AI in software development continues to accelerate, but the real shift underway is less about autonomy and more about addressing the long-standing structural challenges that hold enterprises back. Rather than simply automating tasks, AI is beginning to be leveraged to rebuild the core of software delivery: modernizing legacy systems, improving architectural integrity, strengthening quality and stabilizing pipelines.

AI-first software delivery (AIFSD) represents the end-to-end integration of generative and agentic systems into the full lifecycle of developing software — requirements, design, development, testing, deployment and maintenance. As these capabilities mature, they don't just speed up delivery; they reinforce the foundations on which delivery depends.

Systems will increasingly learn from product goals, user behavior, telemetry and operational signals, enabling continuous improvement. But these capabilities must sit within strong engineering oversight to avoid compounding technical debt, introducing vulnerabilities or creating brittle architectures. The opportunity is not fully autonomous development — it is AI-enabled core renewal.

While generative AI tools can dramatically accelerate delivery and will change the way developers work, it's important to maintain a balance. Regardless of how sophisticated they may become, AI systems must operate under rigorous engineering oversight. Without this, they risk introducing technical debt, security vulnerabilities or hallucinated requirements. Unchecked, AI-generated code may bypass proper architecture practices, or create subtle flaws that later prove costly to fix. Industry analyses have shown that generative AI can lead to maintainability concerns or vulnerabilities if governance, review and validation are not baked into the process.

AIFSD has to be seen as a practice where human engineers and AI systems co-construct software in a complementary way, with the AI handling repetitive, scaffolding and optimization tasks; yet always operating under human-in-the-loop stewardship to ensure accuracy, security and architectural integrity. Enterprises that strike the right balance will develop powerful adaptive capabilities that free them from tech debt and ready them to respond to change.

How AI rebuilds the core

Key trends

Goal-based development environments (GBDEs)
Developers will verbally specify objectives like “build a scalable user onboarding flow,” while AI agents negotiate trade-offs, select libraries and assemble implementations. This shift collapses the distance between business intent and software creation.

Continuous learning delivery systems
Delivery pipelines evolve into neural delivery loops — closed systems where feedback from users, telemetry and market signals immediately inform the next iteration or release.

Neural software twins
Digital twins are already being paired with generative AI in physical systems (see McKinsey, 2024). In software, the idea is nascent, but the analogy holds: maintain a living model of the running system (code + data + performance) that the AI can experiment upon or predict regressions and propose changes before applying them in production.

Synthetic engineers
In research settings, early prototypes of synthetic engineers — composite AI entities composed of multiple specialized models — are beginning to manage entire development streams. These entities collaborate, debate and self-correct to deliver complex systems with minimal human intervention. Though experimental, they hint at the emergence of collective AI design teams.

Multimodal collaboration and real-time translation
As translation, voice synthesis and multimodal reasoning mature, AI plays the role of universal collaborator, instantly bridging gaps between design, engineering, QA and product in different languages and modalities. This ensures global collaboration at human speed.

Signals

AI is already reshaping foundational software engineering practices:

- **AI-augmented development environments** can help streamline testing, debugging and code quality, easing technical debt pressures.
- **Predictive quality engineering.** AI systems are now capable of predicting test coverage gaps, security vulnerabilities and performance issues before code is deployed. This is enabling more proactive approaches to software quality assurance and vulnerability remediation that is providing critical leverage in industries like e-commerce and financial services.
- **AIOps-enhanced pipelines.** The incorporation of AI into DevOps pipelines, or AIOps, is automating incident detection, root-cause analysis and self-healing actions, helping optimize resources and making DevOps teams more productive.
- **Intent-based, goal-oriented coding.** Frameworks are appearing that enable developers to put autonomous coding assistants to work simply by expressing what they want to build, rather than how to build it. It is the AI system that decides the best means of achieving the stated objective. A similar trend is cognitive debugging, where experimental tools use large context windows and neural-symbolic reasoning to detect and explain bugs at the intent level, rather than just sharing their syntax.
- **AI-empowered maintenance and continuous architecture improvement.** In examining existing code bases, predictive agents are able to identify decaying dependencies, outdated frameworks or performance regressions humans may miss. AI-driven refactoring tools can also suggest optimizations to enterprise architecture based on performance and scalability data in real time. The net result is a more efficient tech estate where less time is spent 'keeping the lights on.'
- **The rise of 'adaptive governance'.** AI systems are being leveraged to track compliance, license usage and policy enforcement dynamically, improving the oversight of third-party service usage and ensuring software evolves in a responsible way.

Each of these signals contributes to a stronger, more maintainable core.

The opportunities:

By getting ahead of the curve on this lens, organizations can:

Break through modernization barriers
The difficulty of untangling legacy systems is one of the biggest barriers enterprises and development teams face when building for the future. AI-enabled tools like CodeConcise are proving a massive accelerant to this process, slashing the time and effort needed to reverse engineer code bases and complete the shift to modernized architectures that allow teams to focus on creating value.

Elevate the way teams build
Claims about the ability of AI to speed up the software delivery process are often wildly exaggerated, but we've observed a typical uplift of up to 15%. More significantly, by acting as an engine of customer insights, AI is fueling faster, more intelligent product prototyping and design, and helping teams deliver deeply personalized customer experiences.

Lower operational risks and enhanced reliability
Leveraging AIOps for capabilities like code quality checks and real-time alerting increases the likelihood of catching and addressing software problems and vulnerabilities before they become crippling. Studies have found integrating AIOps into incident management typically improves mean time to detection by over 70%, and reduces mean time to resolution by more than 60%.

Enhance developer satisfaction
AIFSD will redefine what it means to 'ship software.' With AI taking over more mundane or repetitive aspects of the development process, teams will move from builders to curators of intelligence, designing the objectives and constraints within which AI systems will operate. This will free up teams to devote more time to higher-value work that supports product innovation and engagement with the business.

What we've done

One of our clients, a leading automotive manufacturer, had developed a complex system over three decades based on a tech stack for which engineering talent was proving increasingly difficult to find. To update this system and ensure it could continue to deliver value, the company was faced with the task of rewriting the entire code base, which ran to millions of lines. Reverse engineering 10,000 lines of code alone was taking an average time of six weeks. Drawing on the CodeConcise tool, we were able to reduce this cycle by two-thirds — potentially saving the company up to 60,000 person-days, and helping teams approach modernization with renewed confidence.

Trends to watch

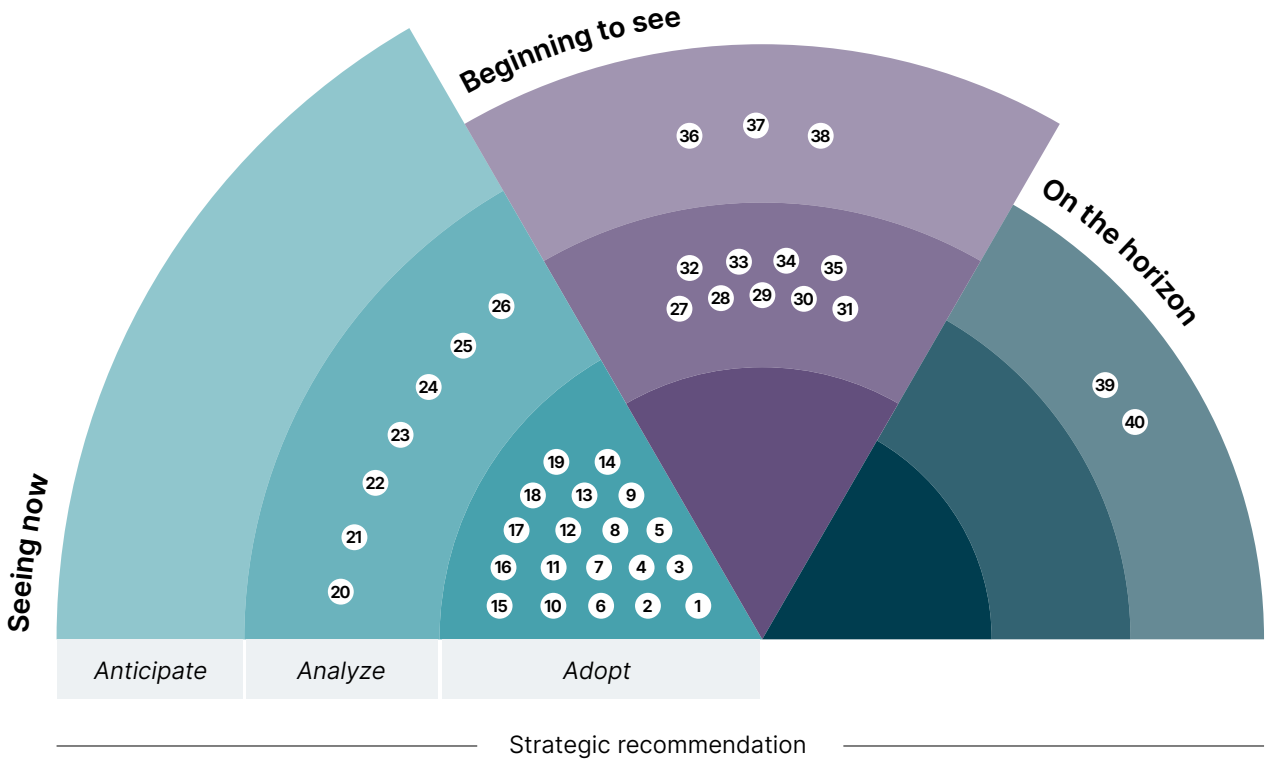

Seeing now

Adopt

1. Generative AI for enterprise
2. Retrieval-augmented generation (RAG)
3. AI code assistants
4. MLOps and model operationalization
5. AI governance and responsible AI
6. AIOps for IT operations
7. AI-assisted coding
8. Software supply chain security/SBOM
9. Developer experience (DevEx)
10. RAG for development context
11. Cloud-native application protection (CNAPP)
12. DataOps and data engineering
13. iPaaS with AI capabilities
14. Low-code workflow automation
15. Unified observability platforms
16. Site reliability engineering (SRE)
17. Configuration management/ policy-as-code
18. Programmable networks/SDN
19. Voice user interfaces

Analyze

20. Cloud development environments
21. Data mesh architecture
22. Data observability
23. Composable enterprise/PBCs
24. API marketplaces and ecosystems
25. Platform engineering for infrastructure and operations
26. Haptic technology

Anticipate - No trends

Beginning to see

Adopt - No trends

Analyze

27. Domain-specific AI models
28. AI for full SDLC
29. Sustainable/Green Cloud Computing
30. Cloud-native application protection
31. AI-augmented engineering teams
32. Software engineering intelligence
33. Data products/ data as a service
34. eBPF kernel-level observability
35. Gesture/motion control

Anticipate

36. Post-quantum cryptography
37. Ambient computing
38. AI-generated UI

On the horizon

Adopt - No trends

Analyze - No trends

Anticipate

39. Artificial general intelligence (AGI)
40. Autonomous AI scientific discovery

Actionable advice

Things to do (Adopt)

- **Grasp the modernization opportunity.** The potential of GenAI to decipher archaic coding languages and address documentation gaps means the time, complexity and cost burdens that have prevented the organization from tackling modernization may no longer apply, easing the path to more resilient, productive systems.
- **Encourage AI adoption, and ensure development teams are 'agent-ready.'** While AI is already a daily reality for the majority of developers, research points to a rise in cynicism about the technology, rooted in concerns about its impact on jobs and issues with accuracy. It's important that enterprises are transparent about their vision for AI; how and where they expect to integrate it into the development lifecycle; and how the use of agents may alter the engineer's role.
- **Institute clear guardrails, checks and balances.** With surveys indicating AI-generated code can introduce serious vulnerabilities, enterprises will need to retool and elevate software quality control and risk management. AI evals are emerging as an essential element of the development process. Teams will also have to view the tools they use, and their outputs, with an ethical lens. Frameworks like the Responsible Tech Playbook present a range of methods and benchmarks designed to ensure development remains principled. A balance has to be struck between instituting clear standards and guidelines, while providing teams room to adopt the technology in the ways that are most relevant for them.

Things to consider (Analyze)

- **Expanding AI's role in the development process.** As systems grow more sophisticated there is clear potential to push use cases well beyond coding into more creative and strategic aspects of development, from evaluating new features to modeling the most likely user response.
- **Rethinking measures of engineering performance.** With human engineers and AI agents increasingly working in tandem, the standard DORA metrics may become less relevant; the speed at which AI can create code vastly exceeds typical deployment benchmarks for example. KPIs will need to be redefined both to account for developers focusing on higher-value tasks, and for systems that will depend on these metrics to make and execute decisions.

Things to watch for (Anticipate)

- **The emergence of neural software twins.** The digital twin techniques often adopted in manufacturing are increasingly pushing into the software domain, allowing for the modeling of entire codebases and system behavior to simulate future states. Such simulations could reduce complexity and introduce more clarity into the development lifecycle, informing future strategies and priorities.
- **The rise of the synthetic engineer.** First came synthetic data; now prototypes of synthetic engineers — composite AI entities composed of multiple specialized models — are beginning to manage entire development streams. These entities are capable of collaborating, debating and self-correcting to deliver complex systems with minimal human intervention. Though in the early stages, they could have major long-term implications for the engineering practice, and hint at the emergence of phenomena like collective AI design teams.

2

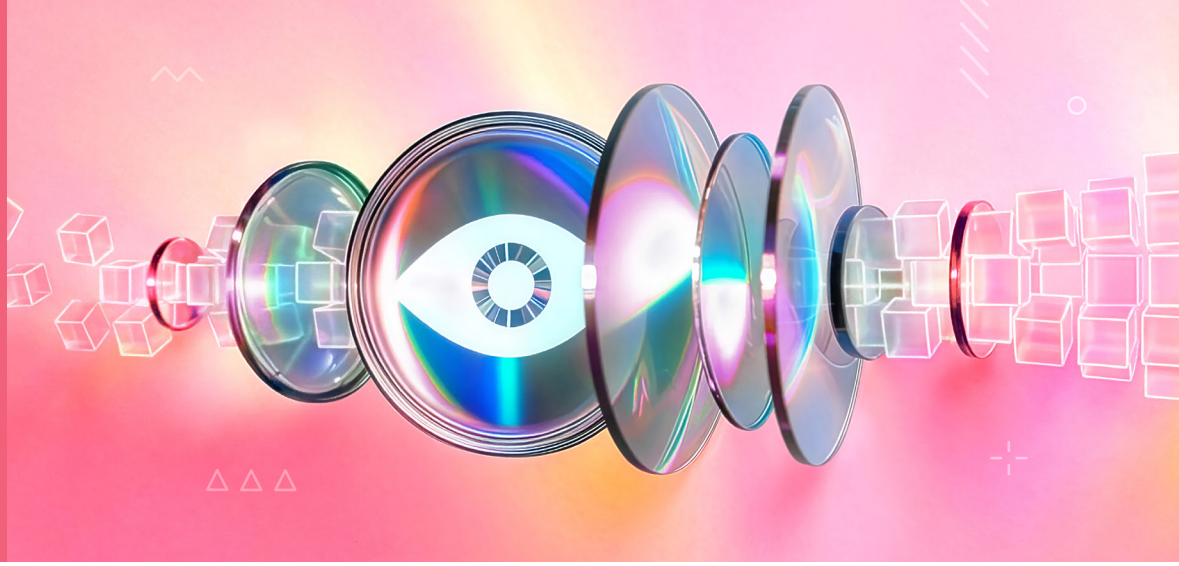

Preparing for agentic transformation

Most enterprises still treat AI as a set of isolated experiments. That mindset is already outdated. In 2026 the question isn't "how many models can we deploy?" but "how fast can we rewire our business so agents can operate across it?" The companies that win won't be the ones building more proofs of concept. They'll be the ones that rebuild their core workflows so intelligence can move freely, act autonomously and deliver outcomes without waiting on human bottlenecks.

The shift is simple but uncomfortable: operationalizing AI is no longer about scaling models. It's about designing architectures where agents can execute work with transparency, guardrails and continuous improvement. If your environment can't support that, you're not ready for AI that works, because you're still in experimentation mode. And as the saying goes, you need to be able to walk before you can run.

When AI moves out of the lab and into production, the organization changes shape. Processes become adaptive. Decision loops compress. Governance becomes embedded instead of documented. And human roles shift from performing tasks to overseeing the behavior of intelligent systems.

This is a profound cultural shift for many enterprises — and such step changes require planning and energy to be done successfully.

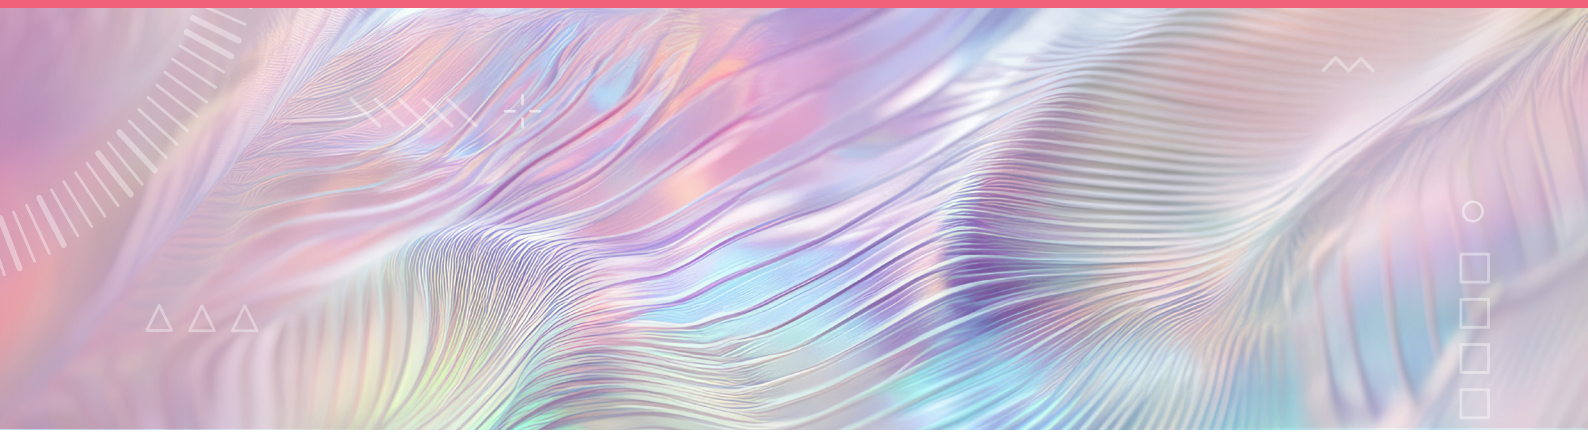

Key trends

Agentic workflows and autonomous operations
AI agents are managing complex, multi-step business workflows, reducing latency between insight and action, and allowing organizations to scale responsiveness and productivity without increasing the burden on humans.

Embedded AI governance
New governance models are integrating AI explainability and transparency directly into operations. This trend matters because trust and accountability will determine the pace of enterprise AI adoption as global regulations tighten.

AI-enhanced decision systems
Businesses are moving from descriptive analytics to adaptive 'co-decision' systems where AI continuously augments human judgment, a major potential boost to agility and precision in fast-moving markets.

Data fabrics and synthetic data ecosystems
Synthetic data and privacy-preserving techniques such as federated learning are unlocking value from sensitive or limited datasets, addressing one of AI's biggest bottlenecks: access to high-quality, compliant data.

Real-time translation and global AI collaboration
Advances from Apple, Google and others in real-time multimodal translation are enabling borderless collaboration across languages and cultures. This will effectively expand the reach of AI-driven businesses, connecting global teams and customers seamlessly.

Signals of this shift include

- **Growth in AI pilots moving to full production.** Some research, like MIT's widely-reported and alarm-inducing 95% study, points to a high proportion of AI projects failing or remaining stuck in the experimentation phase. But these setbacks have to be seen as learning experiences that will inform future successes, and there's equal evidence of businesses extending AI deployments. In one recent poll 58% of enterprises reported AI is now embedded within strategies organization-wide, more than double the previous year's figure.
- **The rise of agentic workflows.** AI agents are now capable of orchestrating multi-step business processes autonomously, from customer support to DevOps. The growing confidence in agents and the expansion of their responsibilities is evident in enterprise implementations of orchestration frameworks like those offered by OpenAI and Anthropic. The overall orchestration market is forecast to nearly triple to over $30 billion by 2030.
- **Integrated AI governance.** As AI highlights the vital importance of data governance to digital leaders, more enterprises are adopting end-to-end governance systems for explainability, compliance and bias detection, transforming responsible AI from policy to platform.
- **Streamlined data-to-value pipelines.** Advances in synthetic data, federated learning and automated labeling are accelerating the data lifecycle from creation to consumption, making AI deployment faster and more secure.
- **The emergence of human-AI co-decision systems.** Businesses are beginning to grasp that AI's real value rests not in automation but its potential to enhance decision-making. More organizations are integrating AI into decision loops where human judgment and machine prediction work together to make informed and impactful choices.
- **The proliferation of AI-native platforms.** Formerly monolithic enterprise systems such as ERP, CRM, HR and finance are evolving into AI-first architectures that embed intelligence into every interaction, boosting operational efficiency and the accuracy of decisions.

Trends to watch

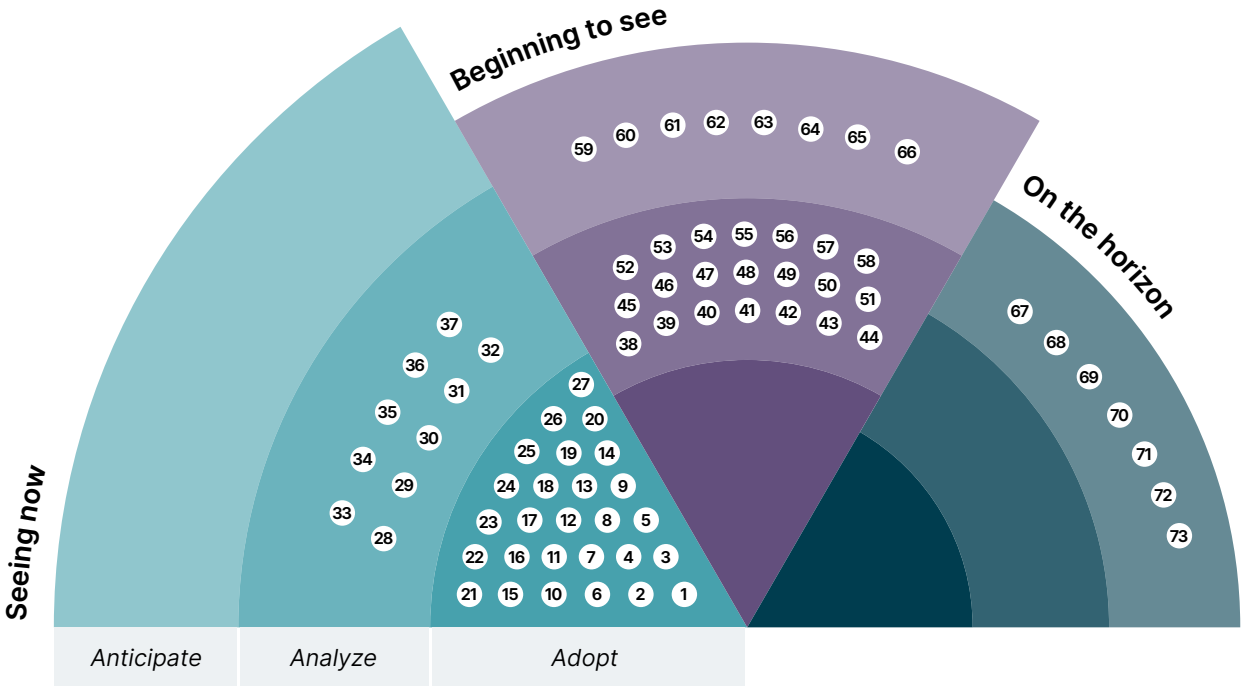

Seeing now

Adopt

1. Generative AI for enterprise
2. Retrieval-augmented generation (RAG)
3. MLOps and model operationalization
4. AI governance and responsible AI
5. AIOps for IT operations
6. iPaaS with AI capabilities
7. Unified observability platforms
8. Configuration management/ policy-as-code
9. Programmable networks/SDN
10. Voice user interfaces
11. AI-ready data infrastructure
12. Computer vision in enterprise
13. FinOps and cloud financial management
14. Cloud security posture management
15. GPU cloud and AI infrastructure
16. Zero trust architecture
17. Extended detection and response (XDR)
18. IAM evolution/ITDR
19. Real-time/streaming analytics
20. Data governance and quality
21. AIOps and GenAI for IT operations
22. Incident management automation
23. Network automation/AIOps
24. AI-powered predictive maintenance
25. IoT security/zero trust for OT
26. RTLS/asset tracking
27. Agentic conversational AI

Analyze

28. Data mesh architecture
29. Composable enterprise/PBCs
30. API marketplaces and ecosystems
31. Distributed cloud infrastructure
32. Event-driven architecture
33. Digital twins for enterprise
34. AI-powered drug discovery
35. Additive manufacturing (3D printing)
36. Private 5G networks
37. Autonomous mobile robots (AMRs)

Anticipate - No trends

Beginning to see

Adopt - No trends

Analyze

38. Domain-specific AI models
39. AI for full SDLC
40. Sustainable/Green Cloud Computing
41. Software engineering intelligence
42. Data products/data as a service
43. Small language models (SLMs)
44. Multimodal AI
45. Industry cloud platforms
46. AI-powered threat detection
47. AI security and governance
48. Continuous threat exposure management
49. Semantic layer/metrics stores
50. AI-embedded enterprise applications
51. Headless commerce/ composable DXP
52. Event-Driven Integration
53. Autonomous vehicles (Level 2+)
54. Intent-based networking
55. LEO satellite connectivity
56. LPWAN/Satellite IoT connectivity
57. Digital Thread in Manufacturing
58. GenAI virtual assistants

Anticipate

59. Post-quantum cryptography
60. AI-powered cloud management
61. AI agents for data and analytics
62. Total experience (TX) platforms
63. Neuromorphic computing
64. Advanced materials (Metamaterials)
65. AI-driven capacity planning
66. Emotion AI/affective computing

On the horizon

Adopt - No trends

Analyze - No trends

Anticipate

67. Artificial general intelligence (AGI)
68. 6G Research
69. General-purpose humanoid robots
70. Machine customers/ autonomous commerce
71. Full autonomous vehicles (Level 5)
72. Programmable money/CBDCs
73. Industrial metaverse

The opportunities

By getting ahead of the curve on this lens, organizations can:

Operate more effectively
While not every effort to scale AI throughout the enterprise proves successful, a substantial and growing body of research points to AI adopters achieving tangible and sustained productivity and efficiency gains.

Choose courses of action with confidence
AI can never be a direct replacement for human judgment, particularly when choices could have far-reaching consequences. But it can play a powerful supporting role by providing data analysis to inform decisions or increase visibility into their possible outcomes. In areas like product pricing AI's contributions to decision-making are directly enhancing revenues and profitability.

Develop more proactive approaches to risk
Embedding AI through operations and processes opens the possibility of assigning agents to constantly monitor enterprise resources, flagging changes in the environment or sudden signs of distress. Leading companies like BMW are using AI-supported systems to identify potential faults before they become bigger problems, avoiding potentially costly disruption and downtime.

Better understand and deliver on customer demands
As AI agents become more integrated into a broader spectrum of customer interactions, the data and insights they generate will generate more granular data, and market and segment-specific insights, enabling marketing and experiences to be tailored with greater precision.

Increase capacity to innovate
AI can augment the enterprise's ability to experiment with and see new concepts through to production both directly — by for example shedding light on areas of opportunity or modeling how a product might function in the field — and indirectly, by allowing team members to carve out more time for creative and collaborative activities.

What we've done

A global pharmaceutical company recognized that traditional engagement models with healthcare professionals (HCPs) were limited by generic tools and static dashboards. In response, the organization partnered with Thoughtworks to rewire its engagement approach using data-driven AI tools that support smarter, more contextual interactions.

At the core of the solution is an AI-powered next-best-action recommendation engine built on a unified data platform that integrates diverse sources and delivers tailored suggestions to go-to-market representatives. Complementing this, a contextual AI assistant provides real-time product insights, content discovery and patient case information to support representative workflows.

These capabilities have boosted effectiveness and trust: representatives access relevant insights faster, address HCP questions with confidence and tailor communications to individual preferences. By embedding intelligence into daily workflows, the organization is transforming engagement from reactive outreach to adaptive, insight-driven interaction, setting the stage for more autonomous support systems in the future.

Actionable advice

Things to do (Adopt)

- **Make AI part of talent strategies.** For AI to be integrated and impactful throughout the enterprise, teams will need to embrace changes in their workflows and understand the broader goals the technology is designed to serve. By some estimates, companies miss out on up to 40% of the potential productivity gains of AI due to skill gaps and anxiety, pointing to the need for ongoing training and support.
- **Define and track AI ROI.** It's also critical to employee buy-in and the wider success of AI initiatives that definitions of value are established and progress measured regularly. Many of the organizations implementing ambitious AI-led transformations still lack clear end-goals or the ability to measure what they achieve.
- **Ensure transparency, and human oversight, in 'co-decision' systems.** As agents take on more initiative within workflows, their reasoning, evidence and data sources must remain inspectable. Techniques like retrieval-augmented generation (RAG) help reinforce explainability so agents operate within verifiable boundaries.

Things to consider (Analyze)

- **Multi-agent systems.** As more companies deploy agents, multi-agent systems, where a number of agents collaborate to perform more complex tasks, are becoming more mainstream and producing positive results in fields like agriculture and supply chain management. Businesses considering agentic AI should ensure they have sufficiently robust data resources to support agentic networks and remain on guard against 'agent-washing.'
- **Tighter AI regulation.** From the US to the UK and India, concerns about cybersecurity and societal impact are fueling plans for tougher rules around the use of AI in multiple contexts. Companies will need to monitor regulatory developments closely and factor a general trend of increased scrutiny and restrictions into their AI plans.

Things to watch for (Anticipate)

- **Adaptive AI ecosystems.** By 2030, enterprises will function as systems of systems where intelligence flows across every process, platform and product. The challenge will move from scaling AI to governing it sustainably.
- **The rise of AI reliability engineering.** With concerns about the (mis)use of AI growing, more leaders will invest in ethical design frameworks and agentic orchestration to ensure resilience and trust. The most successful organizations will not just use AI to optimize decisions; they will reimagine how intelligence itself is operationalized within the enterprise and leverage exemplary AI governance as a differentiator.

3

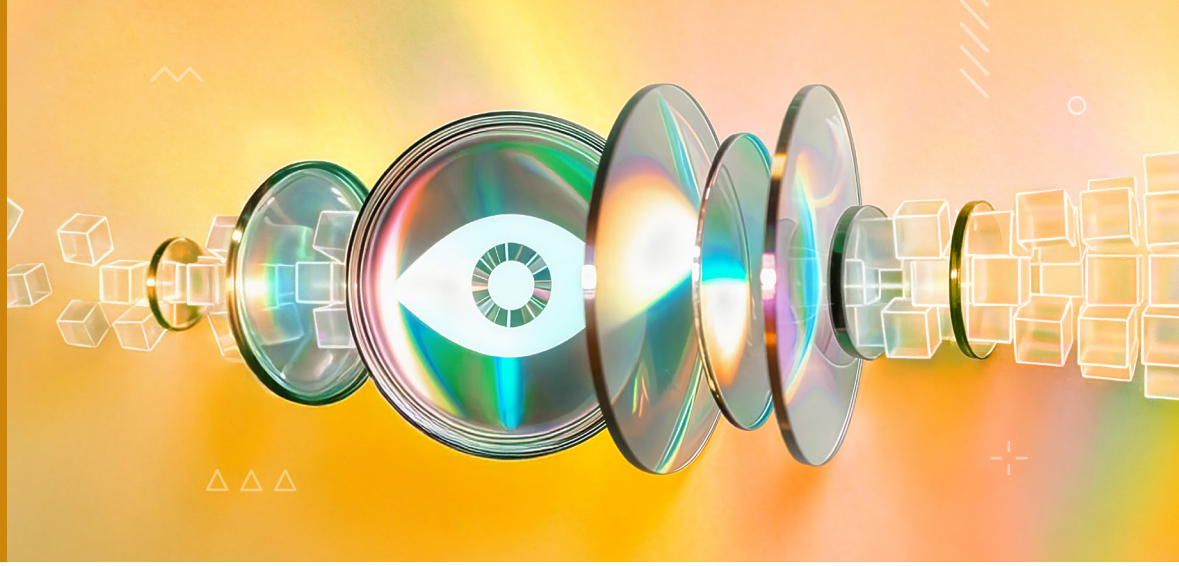

In evolving interactions, AI reimagines the possibilities

Interactions between humans and machines have moved far beyond text, extending into voice, images, gestures and emotional cues. In 2026 and beyond they will move further still, as digital products shift from screen-based interfaces to agentic, intent-driven experiences. Rather than requiring users to master interfaces, systems will increasingly interpret goals, act with autonomy and adapt to context — reimagining how customers access value from technology.

Experiences will be built around agentic interfaces that can take initiative; adaptive systems that sense emotion and environment; and embodied modalities fluent in voice, gestures, gaze and haptics. The emphasis is shifting from designing interfaces to designing relationships between humans, AI agents and the ambient systems around them — a move consistent with the broader shift toward intent-based interaction. Intelligent systems become collaborators that anticipate, learn and co-create outcomes, reducing cognitive friction and expanding creative capacity.

These advances will change how people experience digital systems. Interactions become more natural, contextual and emotionally resonant.

UX evolves into interaction choreography, where designers orchestrate exchanges across multiple intelligences rather than optimizing a single surface.

For enterprises, these shifts redefine the value possible from technology: deeper personalization, richer service ecosystems, more fluid cross-channel experiences and entirely new modes of engagement. They also introduce governance challenges as autonomous systems interpret intent and act independently, heightening the need for transparency and shared control.

Key trends

Real-time translation and cross-cultural interaction
Advances by Apple, Google and OpenAI have made near-instant translation possible in the voice, video, and augmented reality contexts. This trend is crucial because it dissolves language barriers, enabling global collaboration, inclusion and accessibility in real time across both personal and professional interactions.

Agentic UX and predictive collaboration
AI systems will increasingly act as collaborators that anticipate user needs, autonomously completing tasks and negotiating actions. This will redefine productivity tools, transforming them from passive utilities to proactive partners that adapt dynamically to context to achieve the desired outcomes.

Emotionally intelligent systems
With improvements in affective computing and biometric sensing, interactions will become more emotionally aware. Systems capable of recognizing tone, facial expression or physiological signals can respond with empathy to a user's emotional state, which is vital for building trust in autonomous technologies.

Adaptive and contextual interfaces
Experiences will seamlessly shift across devices, environments and modalities depending on where and how users choose to engage with systems. This matters because it creates continuity and consistent personalization, making digital experiences feel like coherent extensions of human intention rather than a grab-bag of disjointed tools.

Collective intelligence and multi-agent collaboration
Enterprises will be able to leverage networks of specialized AI agents that can work alongside human teams to make interactions with customers richer and more responsive, and support faster and more informed decision-making.

Signals of this shift include

- **Rising adoption of agentic systems in customer service.** Tools like OpenAI's agent frameworks and Anthropic's autonomous assistants are extending the role of agents from passive assistants with limited capabilities, to adaptable and active collaborators that can learn from past interactions, reason through and plan the optimal outcomes, and then coordinate to deliver these to the customer.
- **The acceleration of contextual intelligence.** Emotion and context-sensing APIs, combined with wearable data, are allowing systems to respond empathetically to users. The advancement of 'emotionally intelligent' models has reached a stage where they outperform human beings in some psychometric tests.
- **Advances in spatial embodied computing.** Though adoption has often been slower than expected, devices that blend physical and virtual spaces such as the Apple Vision Pro and Meta Quest are growing more sophisticated, user-friendly and, critically, affordable. The buzz around enterprise applications of Apple's Vision Air shows the potential of making interactions more immersive and intuitive in the business context.
- **LLM-native user interfaces.** Language models are now supporting interfaces that can generate, test and modify themselves in real time, based on user intent, and that are being rolled out in products like Google Search and the Gemini app.
- **Collective interaction platforms.** More dynamic, shared dashboards and multimodal dialogue are providing new foundations for humans and AI agents to collaborate, redefining teamwork and coordination.

Trends to watch

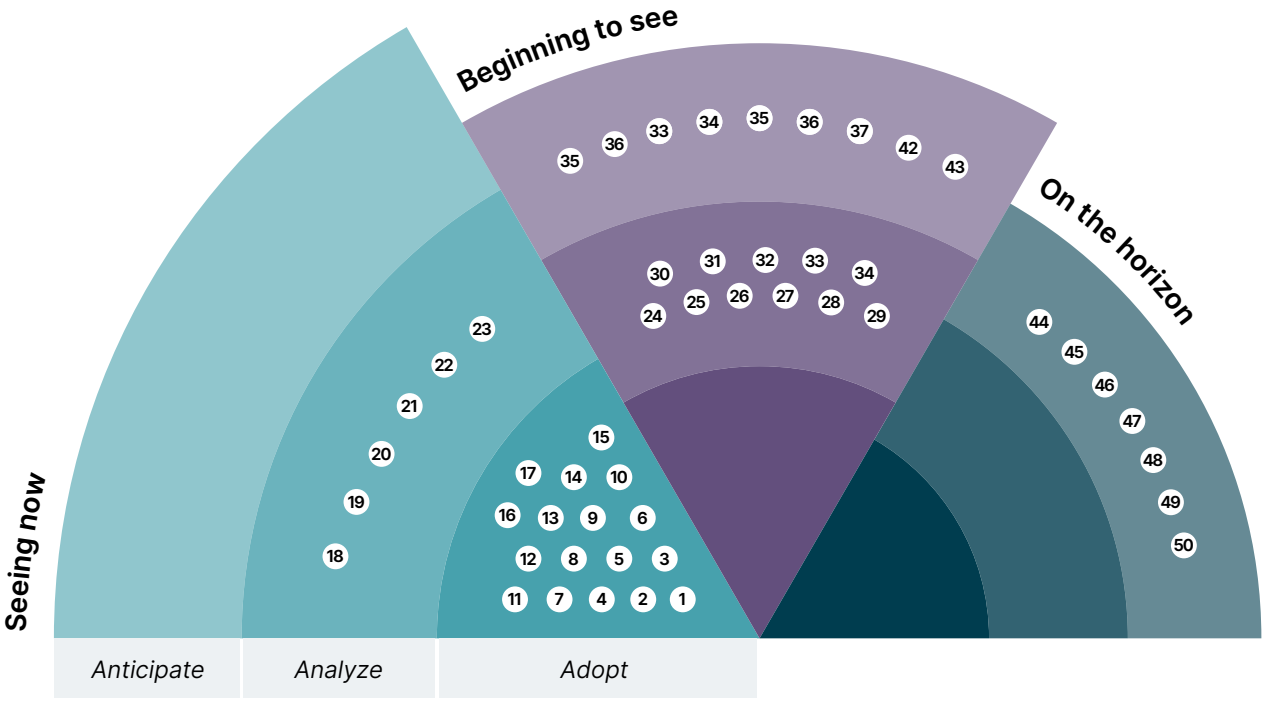

Seeing now

Adopt

1. Generative AI for enterprise
2. AIOps for IT operations
3. Voice user interfaces
4. Computer vision in enterprise
5. AIOps and GenAI for IT operations
6. Incident management automation
7. Network automation/AIOps
8. Agentic conversational AI
9. AI code assistants
10. AI-assisted coding
11. Developer experience (DevEx)
12. Low-code workflow automation
13. NLP for enterprise search
14. Security behavior and culture programs
15. Augmented analytics/ AI-powered BI
16. Data democratization
17. AI-powered accessibility

Analyze

18. Autonomous mobile robots (AMRs)
19. Cloud development environments
20. Platform engineering for infrastructure and operations
21. Haptic technology
22. Passwordless authentication/ passkeys
23. Enterprise AR/VR training

Anticipate - No trends

Beginning to see

Adopt - No trends

Analyze

24. Multimodal AI
25. AI-embedded enterprise applications
26. Headless commerce/ composable DXP
27. Autonomous vehicles (Level 2+)
28. Intent-based networking
29. LEO satellite connectivity
30. GenAI virtual assistants
31. AI-augmented engineering teams
32. Gesture/motion control
33. Spatial computing/ extended reality
34. Connected worker technologies

Anticipate

35. AI-powered cloud management
36. AI agents for data and analytics
37. Total experience (TX) platforms
38. Emotion AI/affective computing
39. Ambient computing
40. AI-generated UI
41. AI video generation
42. Spatial computing/natural UI
43. Enterprise smart glasses

On the horizon

Adopt- No trends

Analyze- No trends

Anticipate

44. Artificial general intelligence (AGI)
45. 6G Research
46. General-purpose humanoid robots
47. Machine customers/ autonomous commerce
48. Full autonomous vehicles (Level 5)
49. Industrial metaverse
50. Brain-computer interfaces

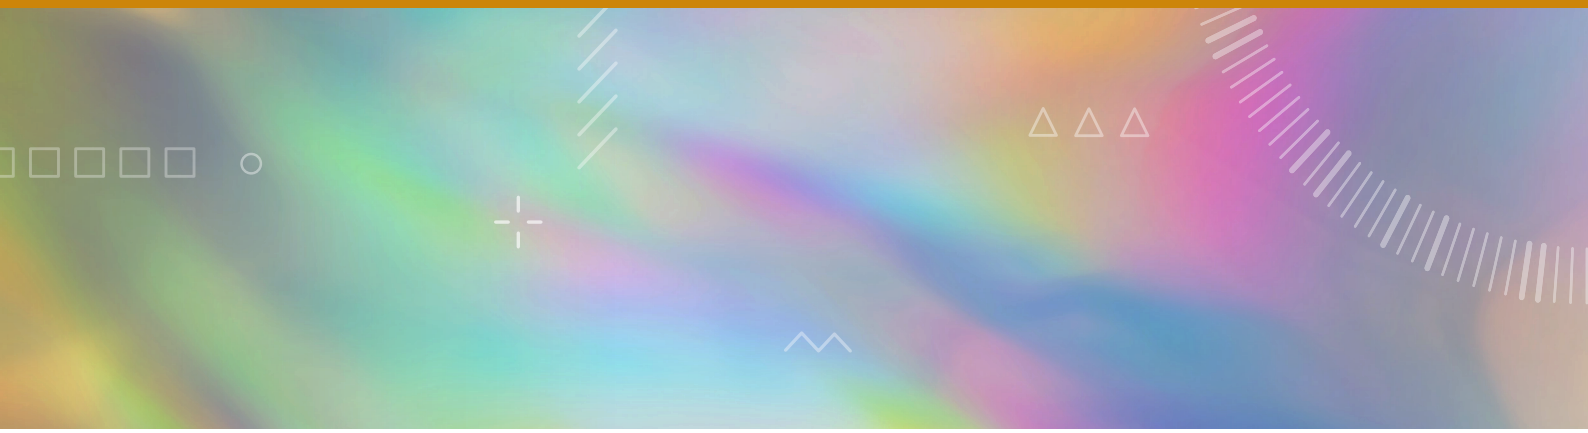

The opportunities

By getting ahead of the curve on this lens, organizations can:

Reimagine the customer experience
One of the main concerns people have about interacting with a system rather than a person is that the exchange will inevitably lack empathy. Affective computing provides a possible means to overcome this barrier and to deliver always-on, responsive and emotionally sensitive customer service at scale.

Further expand the boundaries of accessibility
Advancements in the capacity of systems to recognize and interpret biometric data, gestures and other forms of interaction can be harnessed to level the playing field for collaboration with and improve the delivery of services to groups like the visually or hearing impaired. Similarly, more sophisticated AI interpretation tools have the potential to advance inclusivity for non-native English speakers.

Deliver deeper, and more effective, personalization
Systems capable of recognizing a user's emotional states can learn how to respond to these states in an appropriate and supportive way, adding to the user's sense that their needs are being understood and addressed. Combined with generative interfaces, this opens the possibility of entirely adaptive user journeys that are constantly molding themselves around individual needs, contexts and preferences.

Preempt problems and reinforce the management of risks
Enhanced human-agent collaboration on emerging platforms could have positive implications for practices like cybersecurity, where the always-on monitoring and data analysis delivered by machines should be combined with human oversight. Looking ahead, the integration of VR promises to make human-AI collaboration even more seamless and productive.

What we've done

Reimagining the visitor experience through immersive technology

When the Abraham Lincoln Presidential Library and Museum sought to deepen visitor engagement, it didn't just add another layer of digital content — it reimagined how people connect with history. In partnership with Google Public Sector and Thoughtworks, the museum developed Lincoln Unlocked, an augmented-reality mobile app that blends physical exhibits with immersive digital experiences.

As visitors move through galleries, AR brings artifacts and historical figures to life, unlocking multimedia narratives, interactive tours and contextual storytelling that appeal to diverse audiences. The app also includes accessibility features such as multilingual content and support for visitors with visual or hearing impairments.

By shifting from static displays to dynamic, intent-centered engagement, the museum has expanded its reach, made learning more inclusive and created experiences that resonate across generations. This approach demonstrates how thoughtful use of technology can transform customer interaction, making cultural exploration more vivid, personal and impactful.

Actionable advice

Things to do (Adopt)

- **Think beyond channels to create interaction ecosystems.** Sticking to a standard mix of physical and digital experiences delivered with only slight variations for different devices will mean passing on the chance to forge deeper connections with existing and potentially new customer groups. While it won't be realistic or practical for every enterprise to integrate fully immersive VR or empathy engines into their service offerings, consider where there are clear opportunities to apply evolving technologies to improve the quality of services or enhance accessibility for certain customer segments, and invest accordingly.
- **Design for multimodality.** Customer experiences, journey maps and service flow charts will all need to be updated to factor in a wider spectrum of interactions and channels — such as haptic and voice commands — as these increasingly become table stakes and baked into customer expectations.
- **Establish governance and ethical boundaries.** While customers value interactions with an element of empathy, emotional AI also presents significant risks, from minor misreadings of sentiment to invasions of privacy and out-and-out manipulation. In exploring new means of interaction enterprises should set clear ground rules for how and where emerging technologies should be used, and when there is a need for human involvement.

Things to consider (Analyze)

- **Incorporating AR and VR into the customer experience.** After multiple false starts, the declining price points of AR/VR devices and accelerating adoption in industries like manufacturing could represent a tipping point that will warrant taking the plunge.
- **Exploring the potential of adaptive agentic systems.** Enterprises are moving beyond smaller-scale deployments of GenAI and agents to adopt agentic systems that are capable of taking over a much broader portion of customer interactions. These have significant potential to reduce the burden on human workers, but also require investments in data infrastructure and redefining of the role of service teams to succeed.

Things to watch for (Anticipate)

- **The rise of adaptive environments** that sense and respond seamlessly to customer needs. AI systems will act as 'cognitive exoskeletons' — augmenting human perception, creativity and decision-making. This new frontier will involve designing relationships of trust and negotiated collaboration between humans and intelligent agents. Enterprises that invest now in multimodal design, governance and adaptive AI ecosystems will be best positioned to lead in this next phase of human–machine co-evolution.

4

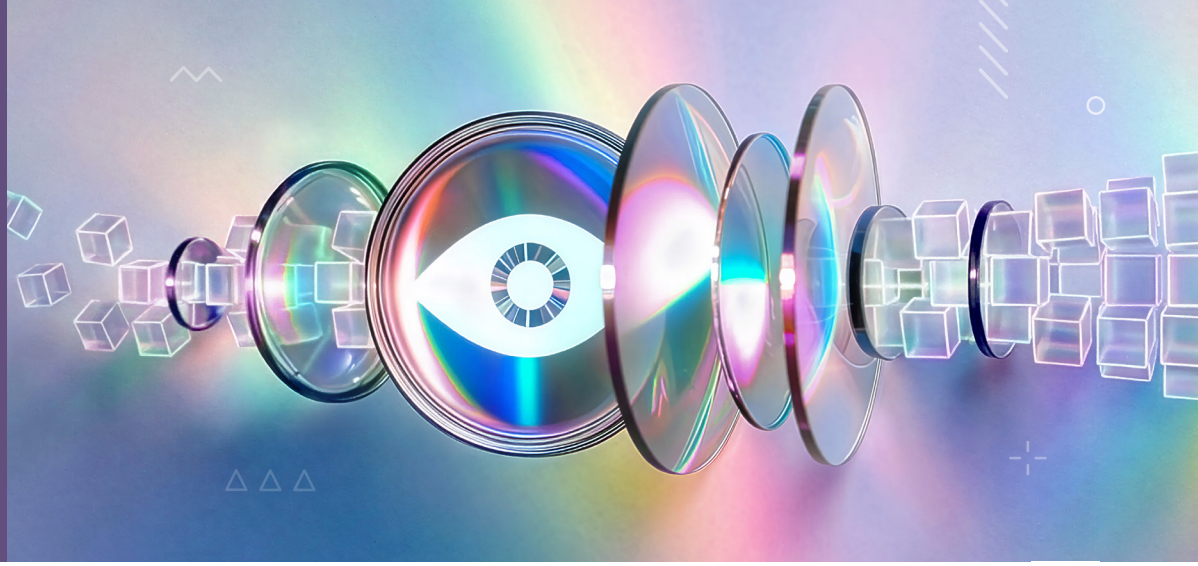

From data platforms to AI-ready data ecosystems

In 2026 and beyond simply having a data platform will not position an enterprise to compete and grow — especially if that platform is a centralized lake that cannot adapt to AI-era demands. As pressure mounts to develop AI-enabled capabilities and products, organizations are learning that the data foundations they built over the past decade are no longer enough. To rebuild the core for the AI era, data ecosystems need to evolve into product-centric, federated environments that can supply trustworthy, real-time data to both humans and intelligent agents.

The goal is no longer a single platform to rule them all, but a dynamic, composable ecosystem ready for composable capabilities that turn modernized data, processes and logic into modular building blocks that teams and agents can reuse, combine and evolve as needs change. This ecosystem becomes an enterprise's enabling layer for agentic systems, grounding them in high-quality data, governed access patterns and clear lineage.

A future-ready data platform encompasses an operating model as well as a technology stack, empowering business domains to create, govern and consume data products on a self-serve basis. Essential emerging features include data product discovery and onboarding, SLAs and quality metrics, lineage-aware ingestion and the golden paths that streamline product development. These foundations position the enterprise to rewire for agents, allowing AI systems to operate safely and at scale.

In this model, data becomes a network of domain-owned products, backed by strong governance and made ready at the edge through real-time processing. Teams can experiment and release AI-powered innovations quickly and securely. Rather than starting from scratch, developers draw on feature and model stores to support training and inference. Evals verify performance and reliability. Safe data access patterns ensure agents can act responsibly across the ecosystem.

A healthy data ecosystem unlocks AI for everything from modernizing the tech estate to reimagining processes and products for agentic workflows. With this network underpinning AI strategy, enterprises gain the speed, insight and productivity that define competitive advantage in the agentic era.

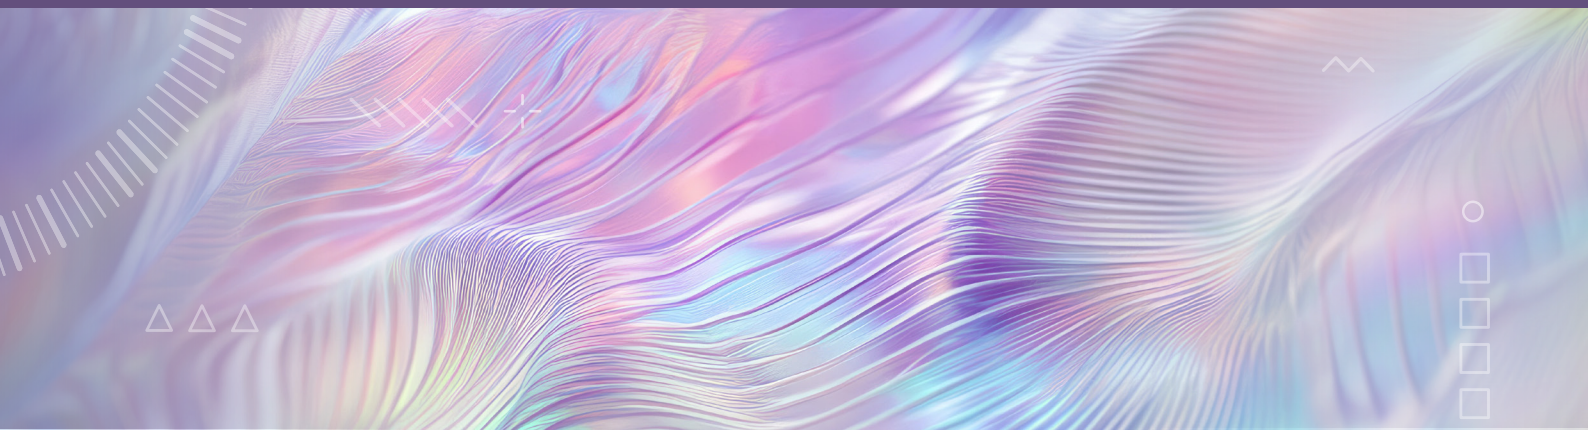

Key trends

Data mesh 2.0: From principles to proven playbooks. Organizations formalize **blueprints** for domain onboarding, data product MLOps and federated governance. Expect standardized templates, data product lifecycle metrics and cost/showback as part of rebuilding core data capabilities.

The composable data product platform. Platform teams assemble interchangeable capabilities (catalog, lineage, quality, transformations, streaming, privacy tooling) behind a coherent developer UX. The emphasis is on golden paths to reduce cognitive load, preparing the environment for agent-ready workflows.

Federated computational governance. Policies (access, PII handling, retention, purpose limitation) are encoded and enforced automatically across pipelines, storage and consumption layers. Governance moves from committees to **real-time controls and auditability.**

Synthetic, privacy-preserving and edge data. Synthetic data augments scarce or sensitive datasets; federated learning and edge processing reduce data movement and strengthen compliance. Platforms standardize the evaluation of utility vs. privacy risk.

AI-ready data value chain. Data platforms integrate feature stores, model registries, eval suites and policy layers. Production AI becomes a first-class consumer of data products, strengthening feedback loops and rewiring for agentic decision making.

Lakehouse and streamhouse interoperability. Teams converge on formats and table standards enabling ACID, time travel and streaming in the same substrate. Operational analytics and real-time ML become simpler and cheaper to operate.

Autonomous data platforms. Platforms evolve toward **self-optimizing systems** that dynamically allocate storage, compute and governance controls based on workload behavior. These platforms learn from telemetry and can reconfigure themselves automatically, reducing maintenance and operational complexity.

Data mesh and agentic AI fusion. Agentic AI systems begin acting as autonomous data stewards — monitoring lineage, evaluating data product quality and suggesting schema or governance improvements. This horizon marks a convergence of AI operations (AIOps) and data governance, leading toward self-healing data ecosystems.

Neural knowledge graphs and semantic infrastructure. Emerging architectures use graph-based embeddings and vector databases to unify structured, semi-structured and unstructured data. These serve as foundations for context-aware, reasoning-capable AI and intelligent retrieval systems across enterprises.

Signals of this shift include

- **A surging appetite for AI development and deployment.** Whether from boards, investors, competitors or their employees, organizations are facing unprecedented demands to integrate AI into workflows, and to develop and launch new AI-powered solutions. Investments are rising accordingly, yet data quality, governance and deciphering the complexity of agentic systems remain formidable challenges. This highlights the need for platforms that embed these capabilities in data and development lifecycles.
- **The diffusion of centralized data platforms.** The limitations of current-generation platforms based on data lakes are becoming more apparent, as they struggle to ingest and harmonize growing amounts of customer data from different sources, or to serve data to users at the speed required for innovation and experimentation. This is driving the evolution of centralized architectures into self-service platforms that empower individual business domains. The shift is not only technical but also organizational; research shows data and data teams being distributed throughout the organization and aligning more closely with different business functions.
- **The maturity of data mesh adoption,** with more enterprises moving from pilots to scaled implementations anchored in domain ownership and data products. The global market for data mesh — an architecture and operating model that decentralizes data resources and accelerates the development of data products — is forecast to nearly quadruple to over $4 billion by 2033 as more enterprises adopt AI-driven data management and real-time analytics.
- **Data governance becoming 'productized' as it rises in importance.** Adoption of data governance programs and policies has surged as awareness of the essential role data quality plays in AI initiatives grows and the high cost of governance failures becomes more apparent. Platforms are changing to incorporate features like policy-as-code, federated computational governance and privacy-by-design patterns as standard.
- **The rise of edge and real-time data processing in platforms.** The proliferation of connected devices and the replacement of traditional batch processing with data streaming means data can be processed faster and ever closer to the source. This has the potential to minimize latency and streamline the processing of workloads, with promising implications for autonomous systems and functions like predictive maintenance.
- **Guardrails emerging around the data stack.** With many leaders lacking confidence in the protection of AI data in their organizations, and a general redoubling of efforts to manage AI risks, more enterprises are looking to integrate evals, guardrails and observability for data and models in data platforms. Such features are also becoming standard issue in the platform solutions being rolled out commercially for capabilities like agentic AI.

Trends to watch

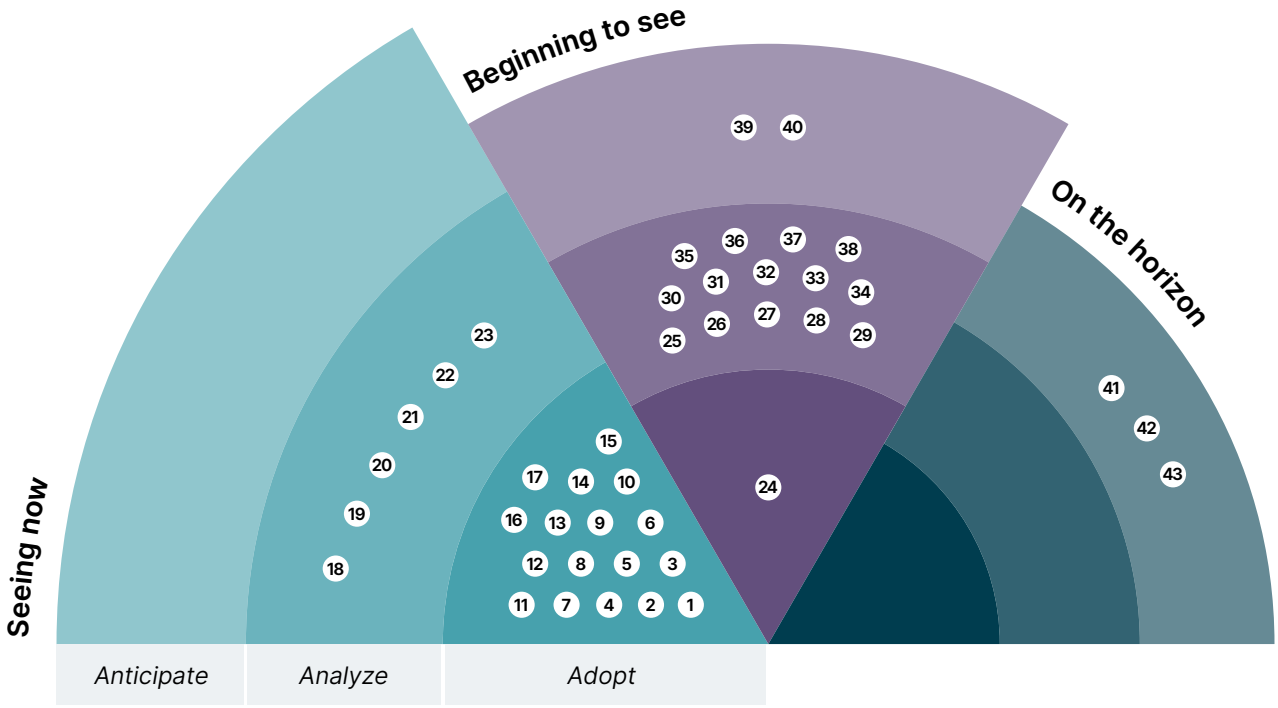

Seeing now

Adopt

1. AIOps and GenAI for IT operations
2. NLP for enterprise search
3. Augmented analytics/ AI-powered BI
4. Data democratization
5. Retrieval-augmented generation (RAG)
6. iPaaS with AI capabilities
7. Unified observability platforms
8. AI-Ready Data Infrastructure
9. GPU cloud and AI infrastructure
10. Extended detection and response (XDR)
11. Real-time/streaming analytics
12. Data governance and quality
13. AI-powered predictive maintenance
14. RTLS/asset tracking
15. RAG for development context
16. DataOps and data engineering
17. Network observability

Analyze

18. Data mesh architecture
19. Distributed cloud infrastructure
20. Event-driven architecture
21. Digital twins for enterprise
22. AI-powered drug discovery
23. Data observability

Anticipate - No trends

Beginning to see

Adopt

24. Data and AI governance convergence

Analyze

25. Multimodal AI
26. Domain-specific AI models
27. AI for full SDLC
28. Software engineering intelligence
29. Data products/ data as a service
30. Industry cloud platforms
31. AI-powered threat detection
32. Continuous threat exposure management
33. Semantic layer/ metrics stores
34. Event-driven integration
35. LPWAN/Satellite IoT connectivity
36. Digital Thread in Manufacturing
37. eBPF kernel-level observability
38. Privacy-enhancing technologies (PETs)

Anticipate

39. AI agents for data and analytics
40. AI-driven capacity planning

On the horizon

Adopt - No trends

Analyze - No trends

Anticipate

41. Industrial metaverse
42. Autonomous AI scientific discovery
43. Synthetic biology for manufacturing

The opportunities

By getting ahead of the curve on this lens, organizations can:

Accelerate time to insight. By smoothing data ingestion, making high-quality data more accessible and facilitating analysis, a modern data platform can significantly increase the speed at which the enterprise is able to identify and act on emerging trends that present opportunities — or risks — for the business. Some of the most prominent examples of AI-powered platforms acting as an accelerant are in drug discovery, where the time required to identify promising compounds is being slashed from months to just minutes.

Reduce regulatory and reputational risks. Low levels of consumer trust and growing regulatory scrutiny mean enterprises using AI in the development process or developing AI products need to act carefully. A platform that ensures data made available to development teams is carefully governed by default, and that activity takes place within defined guardrails, allows the enterprise to innovate in a safe and sustainable way.

Map out future scenarios to make better decisions. AI-ready platforms provide a basis for predictive modeling that can help the enterprise see through likely outcomes before deciding on the best course of action. It's now possible to credibly simulate user behavior and responses to products prior to launch, highlighting areas for improvement or when a product may need to be reimagined altogether. These anticipatory abilities are driving tangible improvements in decision-making, with positive implications for product quality and cost structures.

Enhance resilience. Clear data lineage, consistent data quality signals and the automation of governance reduce the enterprise's exposure to the challenges that accompany wider adoption of AI, from third-party risks to vulnerabilities in AI-generated code.

Shift from a project to a product operating model. The support provided by modern data platforms for data-driven decision making, the rapid testing and rollout of data products and the measurement of outcomes speeds the transition from a project-centric mindset, where products are developed to meet pre-defined requirements, to a product mindset, in which data products are oriented around end-user value. This paves the way for differentiation through customer-centricity and elevated customer experience.

What we've done

A recent engagement with a global medical device maker shows what a modern data platform makes possible in practice. Facing a legacy system that was slow to scale and costly to maintain, the company rebuilt its data foundations around a cloud-first, event-driven architecture. This shift replaced more than 500,000 lines of brittle legacy code with a streamlined, serverless platform capable of handling clinical workloads at volume.

The impact was immediate. Clinical reports that once took hours now run in under ten minutes, even at peak demand. Operational overhead dropped, data quality and consistency improved, and teams gained real-time visibility into device performance and patient outcomes.

These improvements did more than stabilize the platform. They created a data environment that can support more automated decisioning and workflow coordination over time. With data that is timely, trusted and accessible by design, the organization now has the technical footing to explore more adaptive, agent-supported capabilities when they choose to.

Actionable advice

Things to do (Adopt)

- When aiming to build a modern data platform, start with **clearly defined domains and use cases,** gradually extending these to encompass the organization. Avoid platform-first, 'big-bang' variety-builds, which often struggle to scale or deliver value.
- Invest in **developer experience** for data teams, including robust 'golden path' templates, CI/CD, testing and observability. This will ease the path to adoption and position teams to make the most of the platform's capabilities.
- **Treat governance as code and product;** measure it to ensure it is embedded in the platform and the data products being released.
- Align platform metrics to **measurable indicators of business value** such as lead time to new data product, user adoption, reliability SLOs or cost per insight to measure platform performance and make improvements where necessary.

Things to consider (Analyze)

- **Introducing edge processing** to potentially speed up time to insight, minimize the movement of data and strengthen compliance.
- **Experimenting with synthetic data,** which has the potential to enhance functions like market research by bridging the gaps inherent in limited or sensitive data sets.
- **Exploring the fusion of data mesh and agentic AI systems.** As agents are integrated into more enterprise workflows the opportunity will emerge to have them act as autonomous stewards of data, responsible for processes such as tracking data lineage, evaluating data product quality and flagging areas for governance improvement.

Things to watch for (Anticipate)

- Data platforms operating as self-regulating ecosystems, governed by **policy-aware agents** that negotiate access, generate synthetic datasets on demand and spin up ephemeral processing enclaves at the edge.
- More platforms orienting on the North Star of continuous value delivery: generating measurable business outcomes driven by safe, observable data products and powering adaptive agentic AI.
- **Neural knowledge graphs** emerging as a means to enhance flexible governance of modern data platforms, due to their capacity to effectively represent structured data and capture complex data relationships.

5

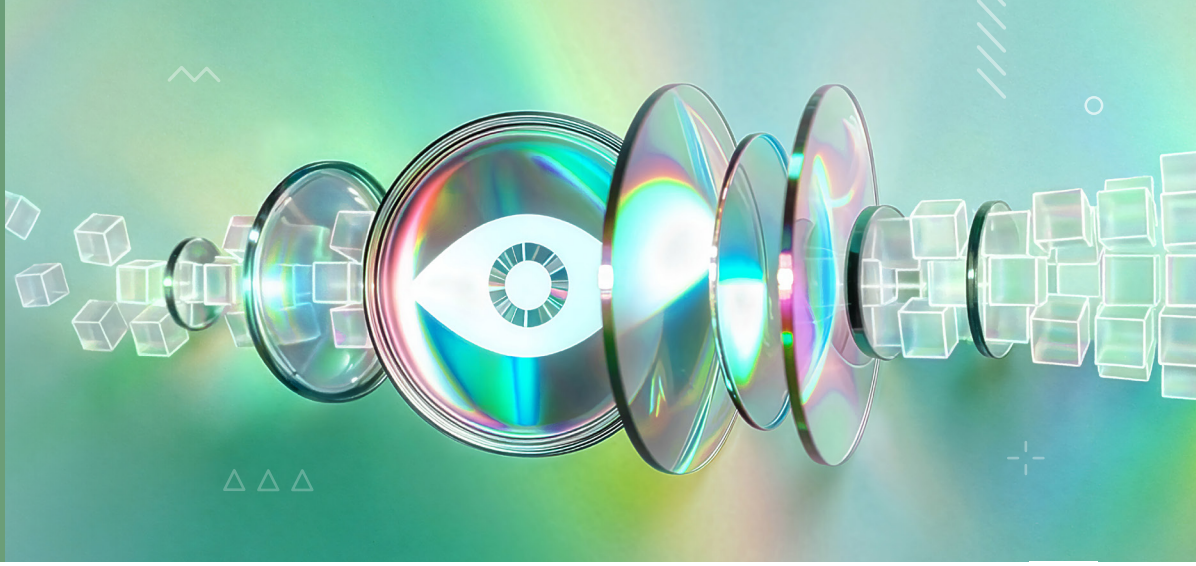

Building your AI future on responsible foundations

As the rapid proliferation of AI raises questions around the ethical use of technology and the stakes of governance lapses, responsible tech is shifting from aspiration to operational discipline. In 2026 and beyond organizations will embed responsibility into technology strategy, architecture and delivery — covering safety, privacy, security, environmental impact, accessibility and social outcomes — not as an afterthought, but as the foundational layer that makes AI systems dependable at scale.

Responsible tech anticipates and manages technology's risks and externalities while maximizing positive impact. In the AI era this spans several connected disciplines:

- **Policy-aware design and engineering.** Privacy-by-design, safety cases, threat modeling, red-teaming, evals and guardrails for generative and agentic AI.
- **Computational governance.** Policy-as-code, automated controls, audit trails and human-in-the-loop escalation.
- **Accountability and transparency.** Impact assessments, provenance, lineage, model cards, incident reporting and user rights.
- **Sustainability and equity.** Resource measurement, inclusive design and community-level harm analysis.

As global and local rules advance, the organizations that succeed will treat responsible tech as governance by design — a built-in capability that underpins every AI-enabled workflow. It becomes the base on which enterprises can adeptly rebuild core systems, rewire for agents and reimagine value through adaptive, intelligent products.

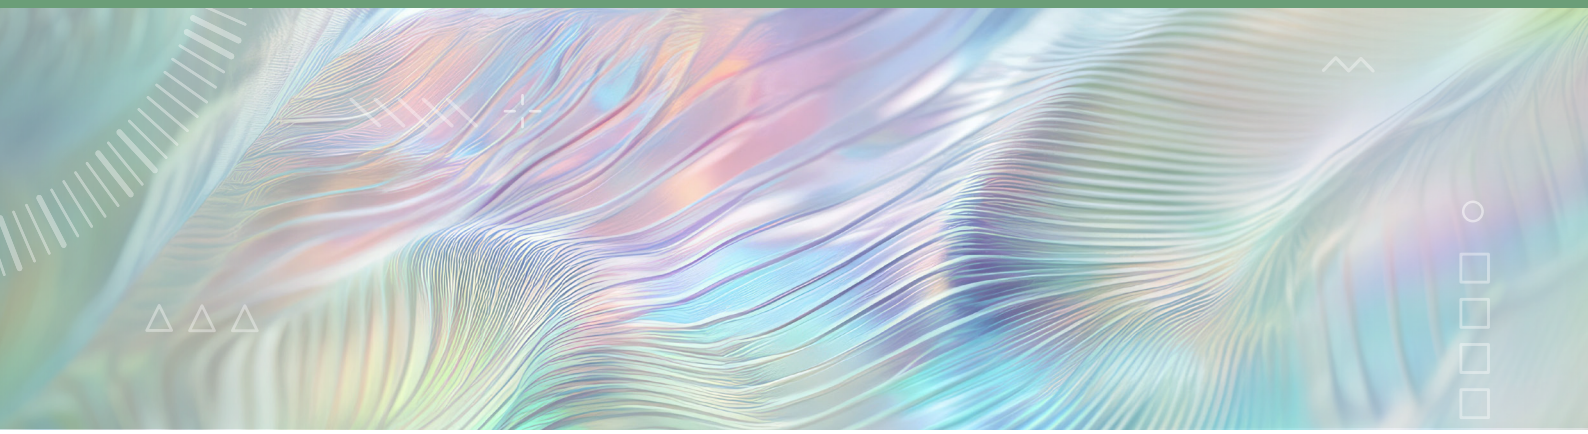

Key trends

Computational governance platforms
Organizations will codify policy (privacy, safety and fairness) as machine-enforced controls integrated with CI/CD, data platforms and runtime. Expect automated DPIAs/AIAIAs, exception workflows and continuous controls monitoring.

Assurance as a product
Evidence packs — traceable datasets, lineage, eval results, risk registers and audit trails — ship with systems. 'Assurance SLOs' become part of service contracts for buyers, regulators and insurers.

Human oversight 2.0
Shift from generic 'human-in-the-loop' to role-specific oversight (safety operator, red-team lead, ethics steward), with well-established escalation trees and shutdown protocols for autonomous agents and decision systems.

AI supply-chain transparency
Model provenance, training-data disclosure windows, third-party component attestations and AI software bills of materials (SBOMs) become table stakes to counter hidden dependencies and license/privacy risks.

Rights-respecting interfaces
UX patterns that operationalize user rights — notice, explanation, contestability, portability — become standardized across sectors and geographies.

Safety cases for agentic systems
Formal safety cases move beyond aviation/medical into AI: structured arguments with supporting evidence for goal-directed agents operating in open-ended environments (e.g., robotics, autonomous operations).

Harmonized global assurance via standards
Convergence of standards such as ISO/IEC 42001 and 23894, and the NIST AI Risk Management Framework enables cross-jurisdictional recognition of controls, reducing the audit burden and enabling portable assurance.

Socio-technical impact markets
Emergence of impact registries and marketplaces for verified positive outcomes (e.g., accessibility uplift, emissions avoided), creating incentives aligned to responsibility metrics.

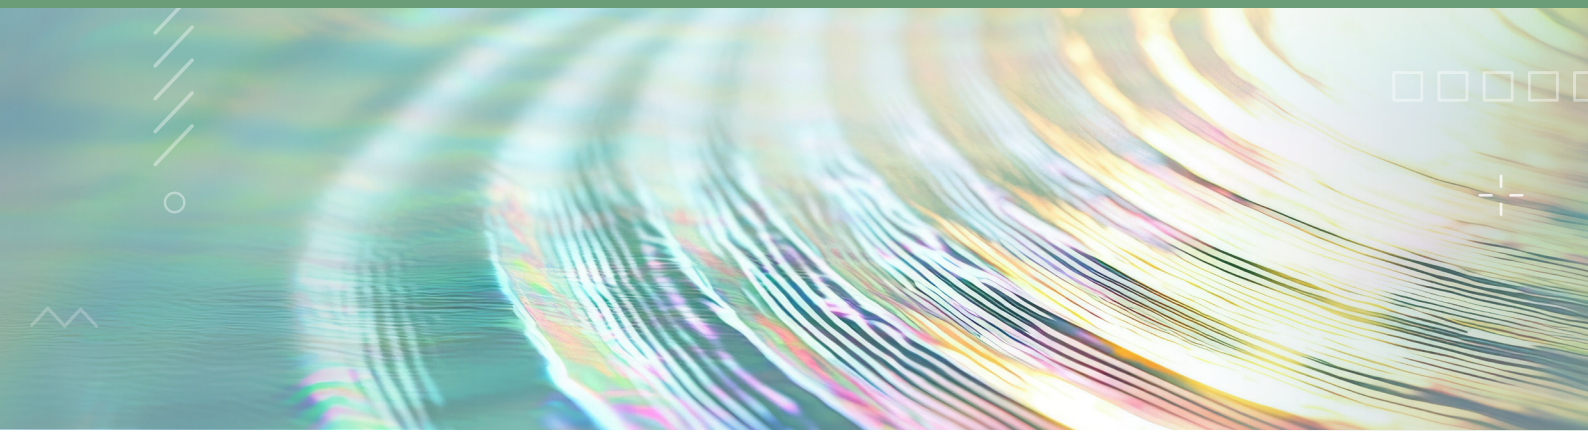

Signals of this shift include

- **The rise of comprehensive, cross-jurisdictional regulation** such as the EU AI Act, which has entered into force with phased obligations. These include bans on unacceptable-risk systems and obligations for providers of general purpose AI models (GPAI) in 2025, as well as stricter rules around the use of AI in high-risk areas like biometrics and law enforcement, originally slated for 2026 but likely to be delayed until 2027.
- **The first internationally legally binding treaty on AI governance,** the Council of Europe's Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, continues to gain momentum, with Canada and Uruguay among the recent signatories.
- **AI rules emerging at the national and state/local level,** though they remain in flux in major economies like the US, where the Biden administration's 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence has been rescinded and followed by new federal directives. Some states are taking matters into their own hands, with California recently passing a law mandating transparency in advanced AI systems; and the implementation of Colorado's landmark AI Act extended to mid-2026.
- **Momentum behind universal standards.** ISO/IEC 42001, which sets requirements for the management of AI systems based on ethical considerations, transparency, and continuous improvement; and ISO/IEC 23894, which provides guidance on integrating risk management into AI-related systems and activities, are examples of the common scaffolding emerging for AI compliance and assurance.
- **Rising corporate appetite, and readiness.** In the business sphere, adoption of AI evals/guardrails, incident response playbooks, and SBOMs for AI supply chains is rising rapidly, as are expectations for transparency and whistleblower protections. More companies are adopting AI codes of conduct and scaling up responsible AI investment plans.

Trends to watch

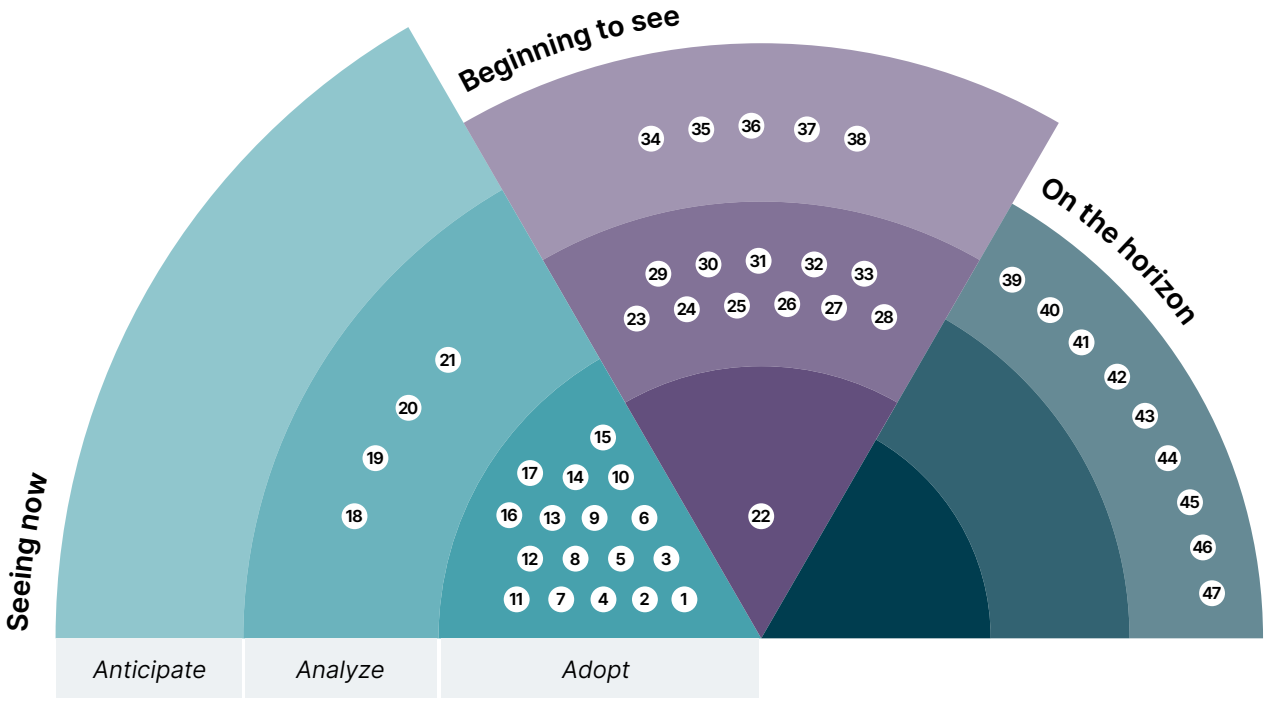

Seeing now

Adopt

1. Extended detection and response (XDR)
2. Data governance and quality
3. Security behavior and culture programs
4. AI-powered accessibility
5. MLOps and model operationalization
6. AI governance and responsible AI
7. Configuration management/ policy-as-code
8. FinOps and cloud financial management
9. Cloud security posture management
10. Zero trust architecture
11. IAM evolution / ITDR
12. IoT security/zero trust for OT
13. Software supply chain security/SBOM
14. Cloud-native application protection (CNAPP)
15. Site reliability engineering (SRE)
16. Third-party risk management
17. Sustainable technology/ green IT

Analyze

18. Distributed cloud infrastructure
19. Data observability
20. Passwordless authentication/ passkeys
21. Grid-scale energy storage

Anticipate - No trends

Beginning to see

Adopt

22. Data and AI governance convergence

Analyze

23. Industry cloud platforms
24. AI-powered threat detection
25. Continuous threat exposure management
26. eBPF kernel-level observability
27. Privacy-enhancing technologies (PETs)
28. Autonomous vehicles (Level 2+)
29. Connected worker technologies
30. Sustainable/ Green Cloud Computing
31. Small language models (SLMs)
32. AI security and governance
33. Cloud-native application protection

Anticipate

34. Emotion AI/affective computing
35. Ambient computing
36. AI video generation
37. Post-quantum cryptography
38. Neuromorphic computing

On the horizon

Adopt - No trends

Analyze - No trends

Anticipate

39. Synthetic biology for manufacturing
40. Artificial general intelligence (AGI)
41. Full autonomous vehicles (Level 5)
42. Brain-computer interfaces
43. Programmable money/CBDCs
44. Commercial nuclear fusion power
45. Space-based solar power (SBSP)
46. Solid-state batteries at scale
47. Quantum internet/ quantum networking

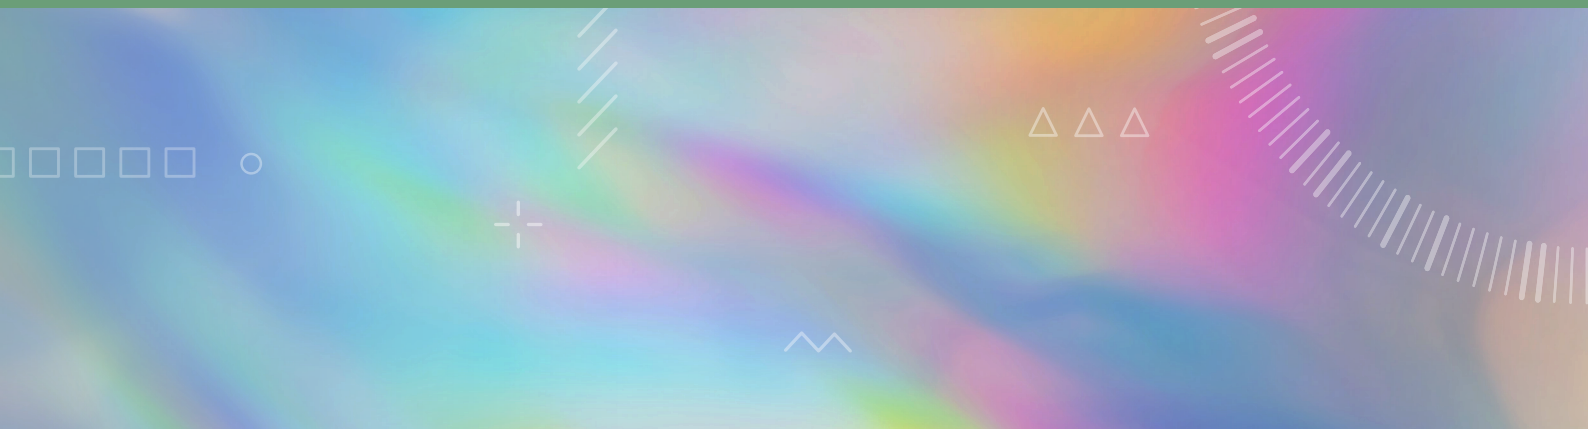

The opportunities

By getting ahead of the curve on this lens, organizations can:

Reduce current and future regulatory exposure
With Gartner predicting that by 2026 around half of governments worldwide will mandate the use of responsible AI through a combination of rules, policies and data privacy requirements, organizations that embrace rigorous governance standards early will be ahead of the curve.

Reinforce customer and stakeholder trust
As the use of AI in areas like marketing and customer service grows, research points to a split emerging, with similar proportions of consumers growing more and less confident that AI tools will handle their data responsibly. Customers are also clear that responsible data and AI practices are an important consideration in their purchasing decisions. This points to a critical juncture where businesses able to convince consumers that they stand on the right side of privacy and ethical lines will gain more market share.

Boost resilience
AI introduces numerous new threats and vulnerabilities, from more sophisticated phishing attempts, to 'back door' breaches via third-party vendors and weaknesses in automated code. Clearly defined policies and practices around data protection and the use of AI systems can reduce the likelihood of major incidents, and limit the fallout when they do occur. In one recent poll over half of business leaders reported responsible AI had improved their cybersecurity and data protection capabilities.

Deliver innovation faster, and with greater confidence
Embedding responsibility as a design input and platform capability, rather than a project phase or box-ticking exercise, ensures products require less in the way of compliance checks and fixes, and are more robust, when they are launched into the market.

What we've done

Thoughtworks partnered with a leading pharmaceutical company on a generative AI chatbot designed for the high-stakes world of preclinical drug discovery. Beyond technical innovation, the project was engineered for responsible use in a scientific setting where trust is non-negotiable. The chatbot combines retrieval-augmented generation and text-to-SQL within an intelligent multi-agent system to unlock insights from large volumes of structured and unstructured preclinical data.

Crucially, the solution prioritizes explainability and reliability through robust engineering practices such as granular citations, error handling, state persistence and LLM fallbacks. These safeguards help researchers understand how answers are formed, trace claims back to source material, and reduce the risk of hallucination-driven decision making.

By enabling fast search, synthesis and report generation across thousands of historical study pages, the chatbot improves speed and accuracy while reducing costly rework.

Actionable advice

Things to do (Adopt)

- **Develop, articulate and model actionable AI policies.** These should include guidelines that will be clear even to a non-technical audience on when and how it's appropriate to use GenAI tools; mandatory steps to establish a baseline of data security; and guidance ensuring AI usage preserves data privacy and intellectual property rights. Policies should be accompanied by regular training so employees are familiar with shifting AI risks.
- **Enhance oversight by strengthening leadership AI literacy.** Senior executives and the board should be kept abreast of developments, threats and opportunities in the AI space and how these could impact their functions and responsibilities, as well as the organization's overall strategy. Ensuring leaders have access to a steady stream of insights and expertise will help them make the right decisions as AI integration progresses throughout the business.
- **Factor AI risks into third-party relationships, as well as the contracts that govern them.** Current and potential vendors should be assessed on a regular basis to ensure their technology practices align with the enterprise's ethical and governance standards.

Things to consider (Analyze)

- **Benchmarking against the best emerging standards.** While some of the rules and standards emerging globally will not be strictly mandatory or may face delays in implementation, they can still be used as a reference point to refine the organization's approaches to responsible technology, and claim competitive advantage by exceeding customer and regulatory expectations.
- **Making responsible tech a platform feature.** Structures like data mesh are making it more feasible to deploy policies and governance across systems. Tools are emerging that will help organizations reflect responsible tech policies in machine-enforced controls that are seamlessly integrated into existing platforms, reducing reliance on human compliance.

Things to watch for (Anticipate)

- **Responsible tech operating as an 'always-on' assurance fabric.** By 2030 responsible tech will be part of the way systems function, showing up in continuous risk sensing, automated controls, periodic third-party attestations, and transparent user rights spanning data, models, and agents.
- **More proactive, AI-enabled risk management.** Business leaders will be able to use increasingly sophisticated policy simulators and behavioral sandboxes to test the socio-technical impacts of solutions before deployment. Tools like these will also support the emergence of impact registries and marketplaces for verified positive outcomes much like those around carbon emissions today, creating incentives aligned to responsible tech metrics.

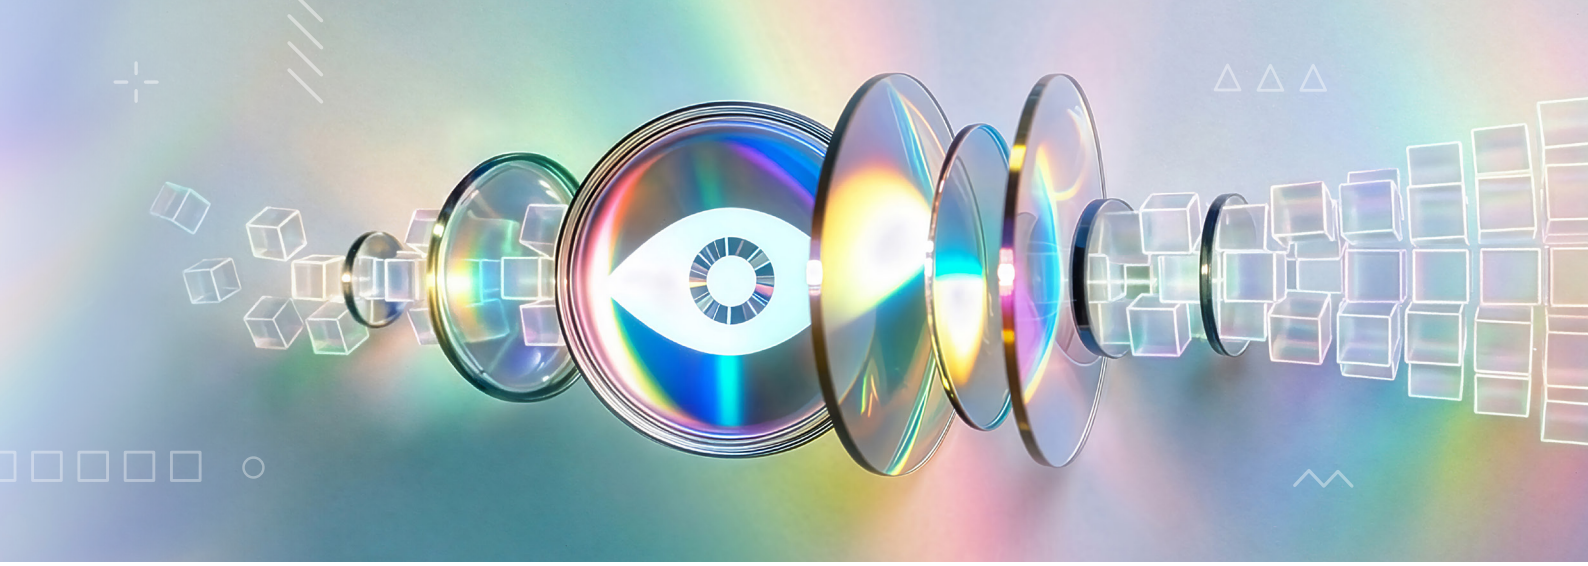

Glossary

A

Additive manufacturing (3D printing): Production using 3D printing for prototyping, custom parts and distributed manufacturing.

Advanced materials (Metamaterials): Engineered materials with properties not found in nature, enabling new applications.

Agentic conversational AI: Advanced chatbots executing multi-step tasks autonomously.

AI agents for data and analytics: Autonomous AI systems performing data analysis, generating insights and executing data tasks.

AI code assistants: AI-powered tools like GitHub Copilot that suggest code completions, generate functions and assist developers in real-time.

AI for full SDLC: AI assistance across the entire software development lifecycle: requirements, design, coding, testing and deployment.

AI governance and responsible AI: Frameworks and policies for ethical AI use, including bias detection, explainability and compliance with regulations like the EU AI Act.

AI security and governance: Securing AI systems and governing their use, including prompt injection prevention and model security.

AI video generation: AI systems creating synthetic video content from text prompts, enabling automated video production at scale.

AI-assisted coding: AI tools providing code suggestions, completions and generation.

AI-augmented engineering teams: Development teams systematically using AI tools throughout the engineering workflow.

AI-driven capacity planning: Using AI to predict infrastructure needs and optimize resource allocation.

AI-embedded enterprise applications: Enterprise software with built-in AI capabilities; 40% of apps will have AI agents by 2026.

AI-generated UI: User interfaces automatically generated and adapted by AI based on context and user needs.

AI-powered accessibility: AI tools improving accessibility through automatic captioning, screen readers and adaptive interfaces.

AI-powered cloud management: Using AI/ML to automate cloud resource optimization, cost management and performance tuning.

AI-powered drug discovery: AI accelerating pharmaceutical development, potentially halving timelines.

AI-powered predictive maintenance: Using IoT sensors and AI to predict equipment failures. 95% report positive ROI.

AI-powered threat detection: Using ML/AI to identify threats, detect anomalies and automate security operations.

AI-ready data infrastructure: Data architecture optimized for AI workloads, including data quality, lineage and accessibility for model training.

AIOps for IT operations: AI-powered IT operations using ML for anomaly detection, alert correlation and the automated remediation of IT issues.

Ambient computing: Technology embedded in environments that responds contextually without explicit interaction.

API marketplaces and ecosystems: Platforms for discovering, consuming and monetizing APIs both internally and externally.

Artificial general intelligence (AGI): AI systems matching human-level intelligence across all cognitive tasks, capable of reasoning, learning and adapting without task-specific programming.

Augmented analytics/AI-powered BI: Analytics tools using AI/ML to automate insight generation, pattern detection and natural language queries.

Autonomous AI scientific discovery: AI systems capable of independently conducting scientific research, generating hypotheses, designing experiments and making Nobel-worthy discoveries.

Autonomous mobile robots (AMRs): Self-navigating robots for material handling and logistics without fixed paths.

Autonomous vehicles (Level 2+): Vehicles with advanced driver assistance, progressing toward higher automation levels.

B

Brain-computer interfaces: Direct communication between brain and external devices. Primarily medical applications currently.

C

Cloud development environments: Browser-based or cloud-hosted IDEs providing consistent, pre-configured development environments.

Cloud security posture management: Tools that continuously monitor cloud configurations for security risks and compliance violations.

Cloud-native application protection: Integrated security platform protecting cloud-native applications across development and runtime environments.

Commercial nuclear fusion power: Fusion reactors generating electricity for the grid, providing nearly limitless clean energy. First commercial plants targeted for mid-2030s.

Composable enterprise/PBCs: Building applications from packaged business capabilities that can be assembled and reassembled.

Computer vision in enterprise: AI systems that analyze images and video for purposes including quality inspection, security and document processing.

Configuration management/policy-as-code: Defining and enforcing infrastructure policies through code for automated compliance.

Connected worker technologies: Wearables and mobile tools enhancing worker safety, productivity and communication.

Continuous threat exposure management: An ongoing assessment and prioritization of vulnerabilities based on real exploitability and business impact.

D

Data and AI governance convergence: Unified governance addressing both traditional data management and AI-specific requirements.

Data democratization: Making data accessible to non-technical users through self-service tools and data literacy programs.

Data governance and quality: Frameworks that ensure data accuracy, consistency, security and compliance. Critical for AI success.

Data mesh architecture: Decentralized data ownership treating data as products owned by domain teams.

Data observability: Monitoring data pipelines for quality, freshness and reliability like infrastructure observability.

Data products/data as a service: Treating datasets as products with defined SLAs, documentation and consumer focus.

DataOps and data engineering: Agile practices for data pipeline development, testing and deployment with continuous integration.

Developer experience (DevEx): Optimizing developer productivity and satisfaction through tooling, processes and culture. More than eight hours are lost every week.

Digital thread in manufacturing: Connected data flow from design through production and service for traceability.

Digital twins for enterprise: Virtual replicas of physical assets with real-time data integration.

Distributed cloud infrastructure: Cloud services distributed to different physical locations while centrally managed, meeting data sovereignty needs.

Domain-specific AI models: AI models fine-tuned for specific industries or functions, delivering higher accuracy than general-purpose models.

E

eBPF kernel-level observability: Programmable Linux kernel instrumentation for deep observability without code changes.

Emotion AI/affective computing: AI detecting and responding to human emotions through facial, voice and behavioral analysis.

Enterprise AR/VR training: Immersive training using augmented and virtual reality.

Enterprise smart glasses: Wearable displays for hands-free information access.

Event-driven architecture: Systems where components communicate through events, enabling loose coupling and real-time responsiveness.

Event-Driven Integration: Integration patterns using events for loose coupling between enterprise applications.

Extended detection and response (XDR): A type of unified security platform that correlates data across endpoints, networks, cloud and email for threat detection.

F

FinOps and cloud financial management: Practices for optimizing cloud spend through visibility, accountability and continuous optimization.

Full autonomous vehicles (Level 5): Vehicles capable of self-driving in any conditions without human oversight, enabling true driverless transportation.

G

GenAI virtual assistants: AI assistants powered by generative models for natural conversation and task completion.

General-purpose humanoid robots: Human-shaped robots capable of diverse physical tasks in unstructured environments, working alongside humans in homes and workplaces.

Generative AI for enterprise: AI systems that create new content (text, code, images) using large language models, enabling automation of knowledge work and content creation.

Gesture/motion control: Interfaces controlled through hand gestures and body movements without physical contact.

GPU cloud and AI infrastructure: Cloud-based GPU resources for AI/ML workloads.

Grid-scale energy storage: Large battery systems enabling renewable energy integration and grid stability.

H

Haptic technology: Touch feedback systems providing tactile sensations in digital interactions.

Headless commerce/composable DXP: Decoupled commerce architecture separating frontend presentation from backend functionality.

I

IAM evolution/ITDR: Advanced identity management with identity threat detection and response for real-time identity-based attacks.

Incident management automation: Automated detection, routing and initial response to IT incidents.

Industrial metaverse: Persistent, photorealistic digital twins of factories, cities and infrastructure enabling real-time simulation, optimization and remote collaboration.

Industry cloud platforms: Cloud solutions tailored for specific industries with pre-built compliance, workflows and integrations.

Intent-based networking: Networks configured through business intent rather than device-level commands.

IoT security/zero trust for OT: Securing IoT and operational technology with zero trust principles.

iPaaS with AI capabilities: Integration platform as a service with AI-powered mapping, transformation and anomaly detection.

J

K

L

LEO satellite connectivity: Low Earth Orbit satellite internet providing broadband with 20-60ms latency globally.

Low-code workflow automation: Visual tools for building automated workflows without extensive coding, accelerating process automation.

LPWAN/Satellite IoT connectivity: Low-power wide-area networks and satellite for remote IoT device connectivity.

M

Machine customers/autonomous commerce: AI agents acting as autonomous economic actors, purchasing goods and services on behalf of humans or organizations without human intervention.

MLOps and model operationalization: Practices for deploying, monitoring and managing ML models in production which ensure reliability and effective governance.

Multimodal AI: AI systems processing multiple data types (text, images, audio and video) simultaneously for richer understanding.

N

Network automation/AIOps: Automating network configuration, monitoring and troubleshooting using AI/ML.

Network observability: Deep visibility into network behavior beyond traditional monitoring for troubleshooting.

Neuromorphic computing: Brain-inspired computing architectures for energy-efficient AI processing.

NLP for enterprise search: Natural language processing that enables semantic search across enterprise content, improving information discovery.

O

P

Passwordless authentication/passkeys: Authentication without passwords using biometrics, hardware keys or cryptographic credentials.

Platform engineering for infrastructure and operations: Applying platform engineering principles to infrastructure and operations teams.

Post-quantum cryptography: Cryptographic algorithms resistant to quantum computer attacks; NIST finalized standards in August 2024.

Privacy-enhancing technologies (PETs): Technologies enabling data use while protecting privacy: differential privacy, homomorphic encryption, etc.

Private 5G networks: Enterprise-operated 5G networks for industrial IoT, campus coverage and specialized applications.

Programmable money/CBDCs: Digital currencies with embedded rules enabling automated compliance, conditional payments and machine-to-machine transactions.

Programmable networks/SDN: Networks controlled through software APIs enabling automation and dynamic configuration.

Q

Quantum internet/quantum networking: Networks using quantum entanglement for unhackable communication and distributed quantum computing, enabling new security paradigms.

R

RAG for development context: Using retrieval-augmented generation to provide AI coding assistants with codebase-specific context.

Real-time/streaming analytics: Processing and analyzing data as it's generated for immediate insights and actions.

Retrieval-augmented generation (RAG): Combines LLMs with information retrieval to ground AI responses in verified data sources, reducing hallucinations by 70-90%.

RTLS/asset tracking: Real-time location systems for tracking assets, inventory and personnel.

S

6G research: Next-generation wireless research targeting 2030 deployment with terahertz frequencies.

Security behavior and culture programs: Human-centric security initiatives changing employee behaviors beyond traditional awareness training.

Semantic layer/metrics stores: An abstraction layer that consistently defines business metrics. Particularly critical for AI accuracy per Gartner.

Site reliability engineering (SRE): An engineering approach to operations focusing on reliability through automation and error budgets.

Small language models (SLMs): Compact AI models under 10B parameters optimized for edge deployment, offering 60-80% cost reduction vs. large models.

Software engineering intelligence: Platforms analyzing development data to measure productivity, identify bottlenecks and optimize processes.

Software supply chain security/SBOM: Securing software dependencies and generating a Software Bill of Materials.

Solid-state batteries at scale: Next-generation batteries using solid electrolytes, offering 50-80% higher energy density, faster charging and improved safety for EVs and devices.

Space-based solar power (SBSP): Satellites collecting solar energy in orbit and transmitting it wirelessly to Earth, providing 24/7 baseload clean power.

Spatial computing/extended reality: Technologies merging physical and digital worlds including AR, VR and mixed reality.

Spatial computing/natural UI: Computing that blends digital content with physical space for natural interaction.

Sustainable technology/green IT: IT practices reducing environmental impact through efficiency, renewable energy and carbon management.

Sustainable/Green Cloud Computing: Cloud practices minimizing environmental impact through efficient resource use, renewable energy and carbon tracking.

Synthetic biology for manufacturing: Engineering biological systems to produce materials, chemicals and medicines at industrial scale, potentially replacing 60% of physical production by 2040.

T

Third-party risk management: Managing security risks from vendors and partners. 35.5% of breaches originate from third parties.

Total experience (TX) platforms: Unified platforms that address customer, employee and user experience holistically.

U

Unified observability platforms: Integrated monitoring across metrics, logs and traces.

V

Voice user interfaces: Voice-controlled systems with 95%+ accuracy for major languages.

W

X

Y

Z

Zero trust architecture: A security model where verification is required for every access request, regardless of location.

We are a global technology consultancy that delivers extraordinary impact by blending design, engineering and AI expertise.

For over 30 years, our culture of innovation and technological excellence has helped clients strengthen their enterprise systems, scale with agility and create seamless digital experiences.

We're dedicated to solving our clients' most critical challenges, combining AI and human ingenuity to turn their ambitious ideas into reality.

thoughtworks.com

/thoughtworks

Design. Engineering. AI.