Open Innovation

Disruption Watch Trends

H1 2027

The tech-led trends shaping the future of business

November 2026



Foreword

New technology rarely changes things on its own; cost, regulation, infrastructure and people's willingness to use it decide when it starts to matter.

A change that makes life easier for customers can also leave them with less control over their data, raising questions about privacy and trust. Businesses face a different problem: a technology can alter who owns the customer relationship before it overturns the business model.

This report looks at the signals behind such shifts, including what could slow or redirect them, although we cannot say which changes will endure. We show what the evidence supports and where uncertainty remains and raise the questions we believe are worth asking now.

“Information is the resolution of uncertainty”

- Claude Shannon, polymath and father of information theory



This report has been created to help businesses consider potential future disruption, and the associated opportunities and risks, that may be realised from emerging tech-led trends.

The report does not make any assumptions of how the future might unfold, rather it poses questions that businesses may find useful when considering potential actions in the face of possible disruption.

Because readiness starts with insight.

At NatWest, Open Innovation isn't just about keeping pace with change. It's about translating what this means for our customers and the bank.

Open Innovation helps NatWest Group thrive in a disruptive, fast-evolving landscape - by learning, responding, and growing.

Through partnering across the bank and with external partners, we're collectively building the future of banking through:

1. Experimentation

Testing ideas before making bigger investments.

2. Minority investments

Making investments in emerging technology companies.

3. Venture build

Creating new companies outside the bank.

4. Community

Helping colleagues learn about new technology and growing relationships with innovation economy.

Find out more

Get in touch to learn more about NatWest Group Open Innovation.

Strategic Investments | NatWest Business

OpenInnovation@natwest.com



Disruption
/dɪsˈrʌp.ʃən/

Noun

The action of preventing something, especially a system, process or event, from continuing as usual or as expected and that can result in positive or adverse outcomes.

This report moves from signals to trends, then to three deeper discussions.

The disruption radar

A visualisation of the most important emerging tech and associated developments, focusing on where there is potential for real-world impact and disruption.

The radar tracks nearly 100 signals, from developments in artificial intelligence (AI) to changes in healthcare. Each signal is an emerging technology or a development linked to one.

See slide 9

Trends

Longer-term tech-led changes that we believe have the potential to create future disruption. These are formed from the application and convergence of disruption signals. We highlight 8 trends in this report:

- *Workflows go agentic*
- *From chatbot to assistant*
- *AI's power problem*
- *Trust under pressure*
- *Robots get a brain*
- *Quantum sets a deadline*
- *Healthspan for the ages*
- *Orbit as infrastructure*

See slides 12-19

Double-clicks

The highlighted trends have the potential to positively or adversely disrupt people, business and/or society. We look more closely at three questions these trends raise for businesses and customers:

- *The economics of autonomy*
- *Agentic commerce takes shape*
- *Inside the AI trust paradox*

See slides 23-34



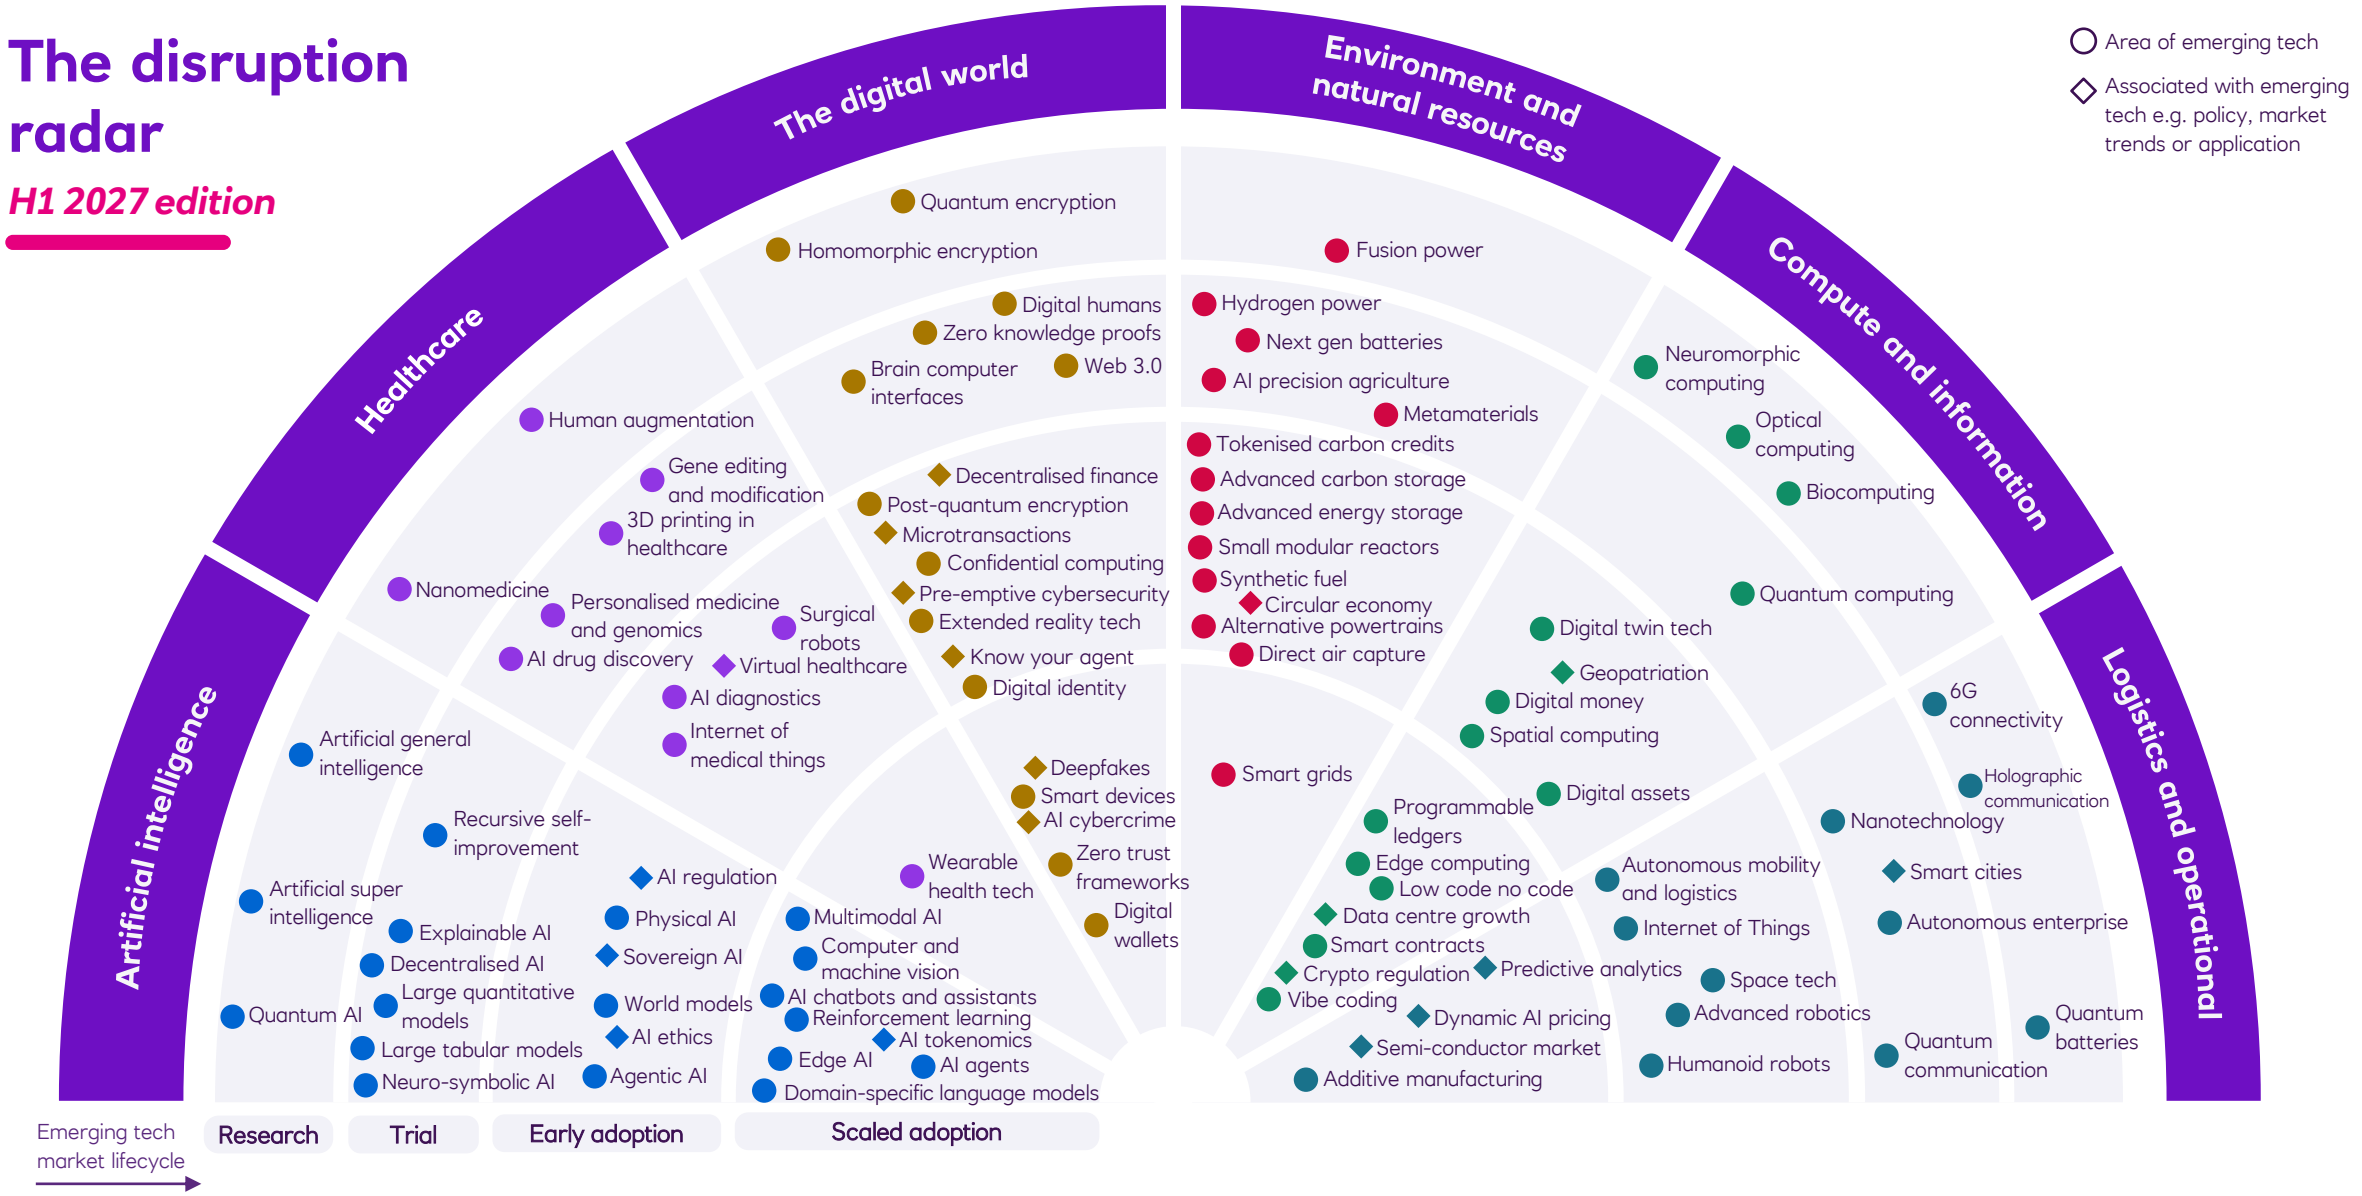
01.

The disruption radar and trends


There is no shortage of new technology. The hard part is deciding what deserves attention now.

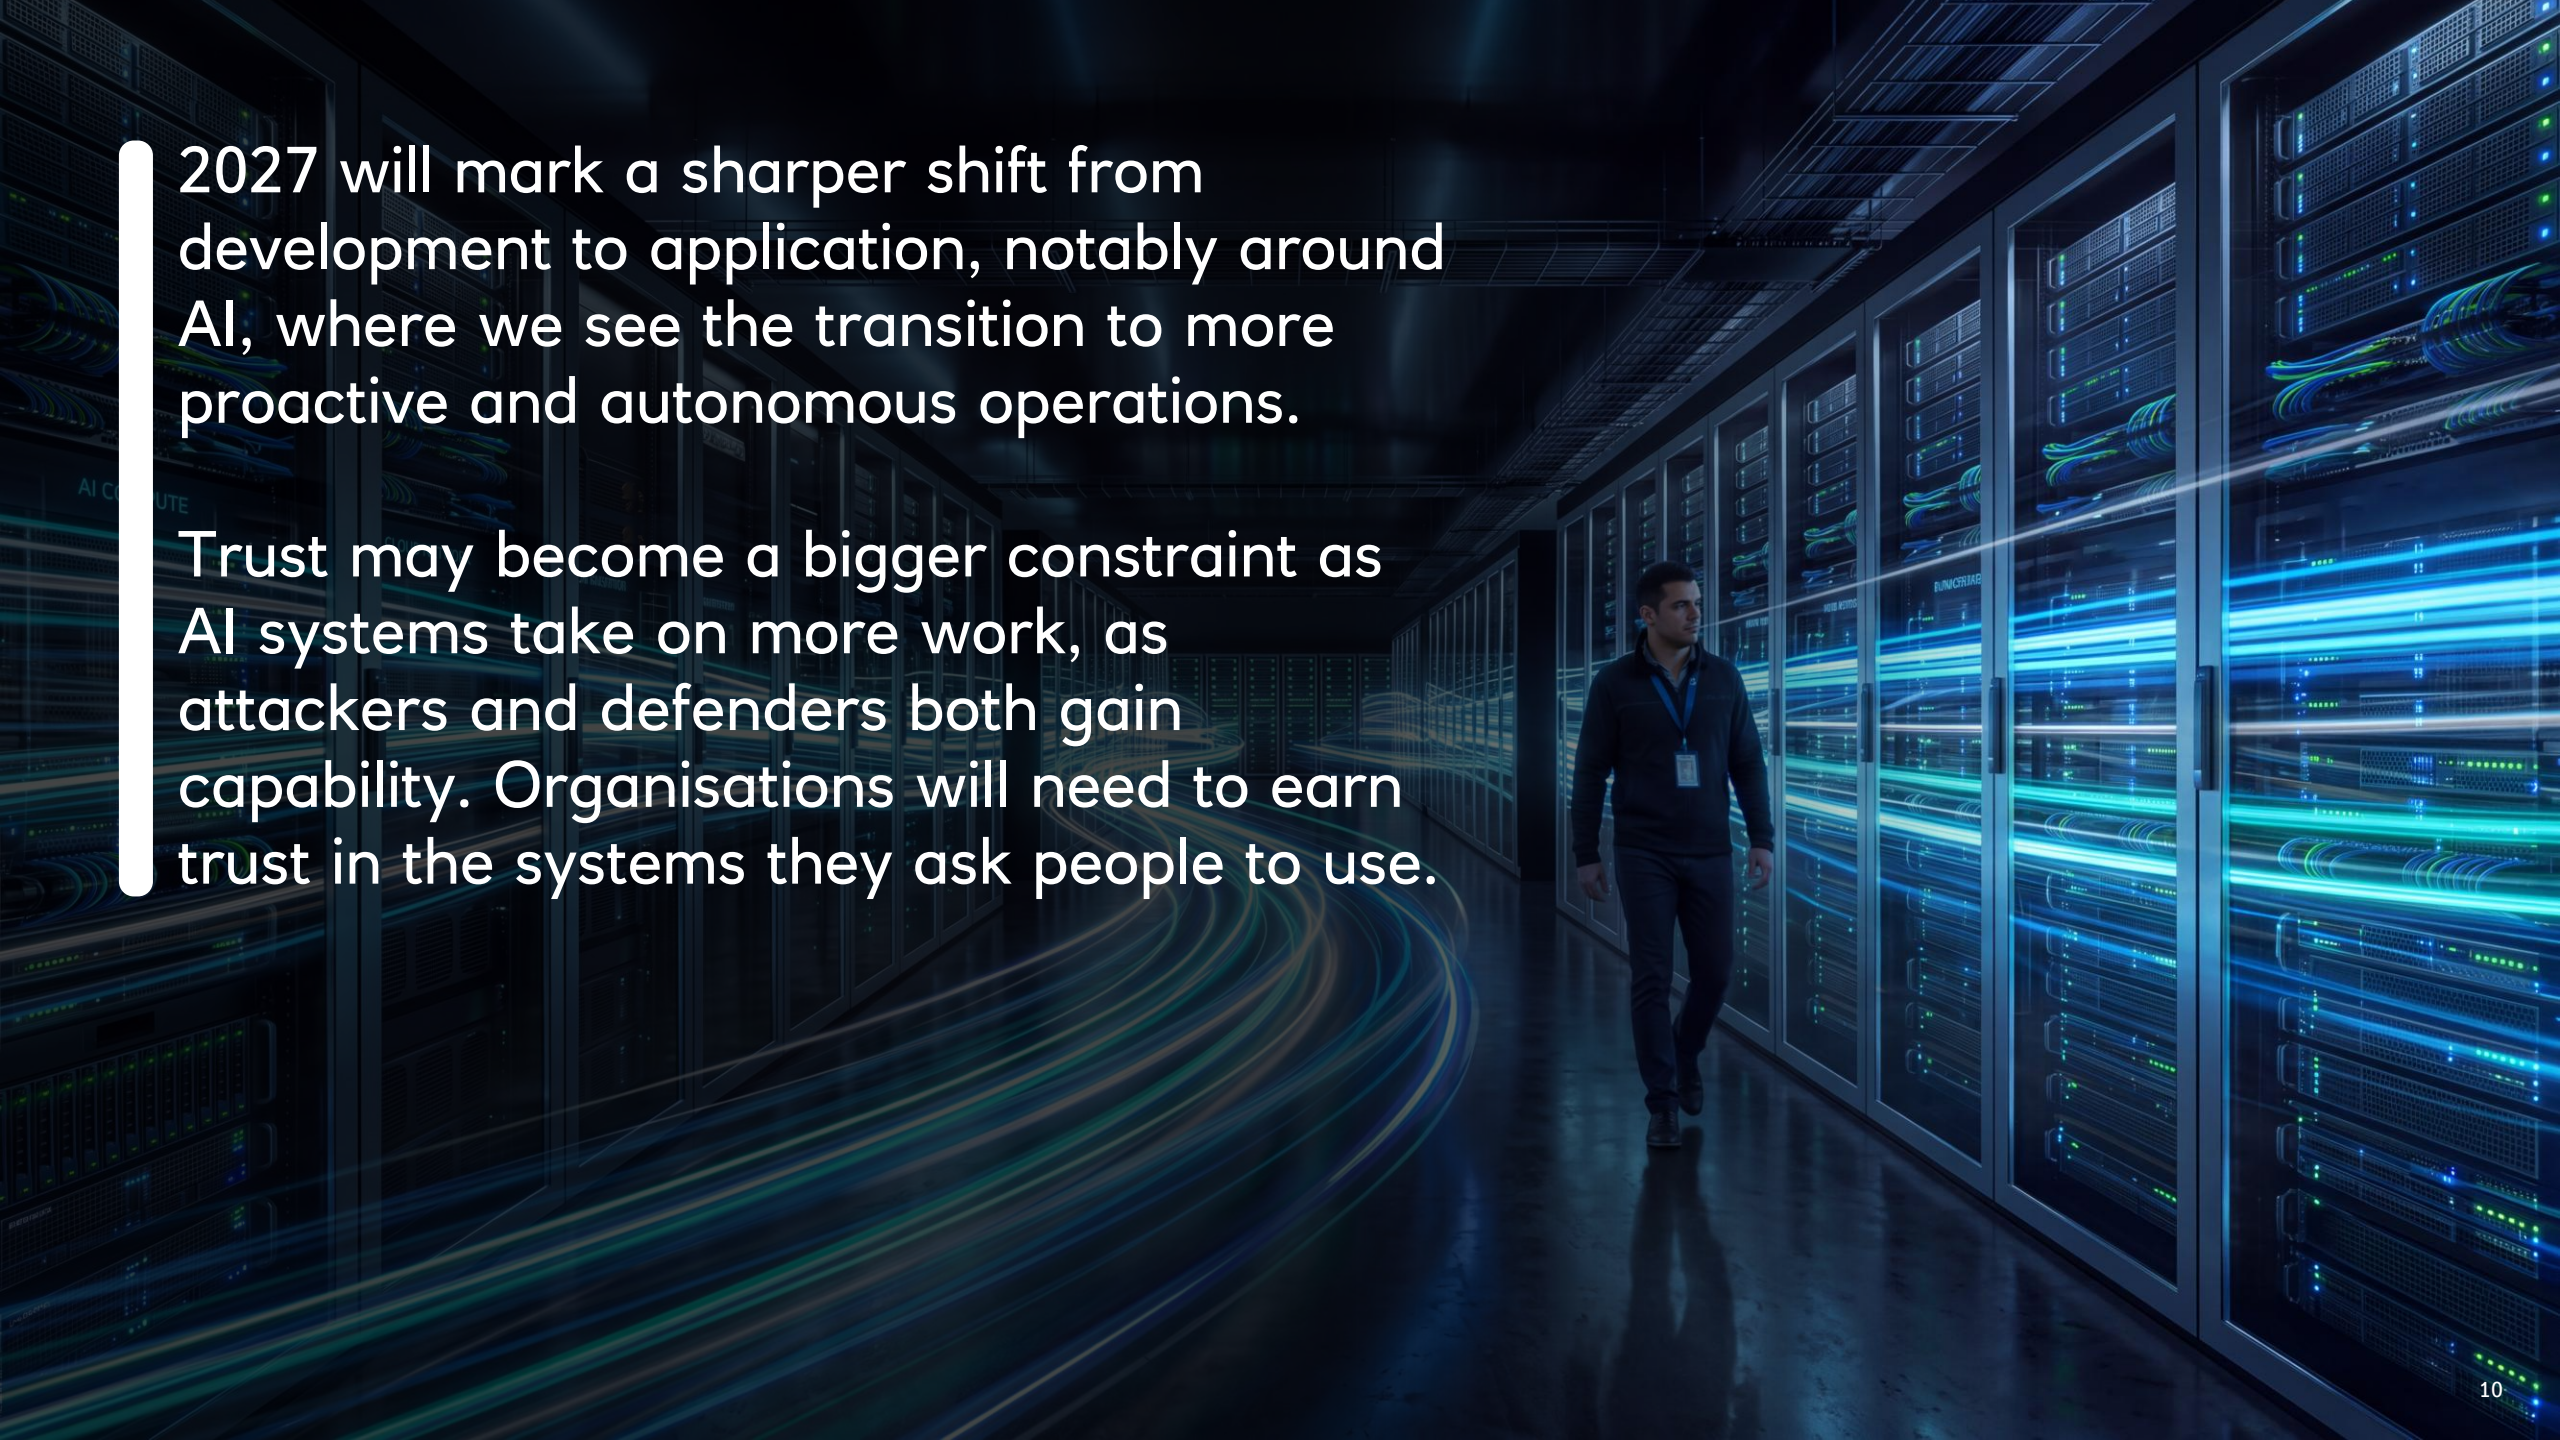
The disruption radar

H1 2027 edition



2027 will mark a sharper shift from development to application, notably around AI, where we see the transition to more proactive and autonomous operations.

Trust may become a bigger constraint as AI systems take on more work, as attackers and defenders both gain capability. Organisations will need to earn trust in the systems they ask people to use.



Eight tech-led trends with the potential to disrupt business

1
Workflows go agentic


2
From chatbot to assistant


3
AI's power problem


4
Trust under pressure


5
Robots get a brain


6
Quantum sets a deadline


7
Healthspan for the ages


8
Orbit as infrastructure


Trend 1

Workflows go agentic

In AI, 'agentic' describes a system that can work towards a goal with limited supervision. It decides what to do next and carries out the work. As models have become more capable, enterprises have started to put agents to work. Some handle customer support or complete payments, while others coordinate operational work such as inventory management.

As foundation models improve, platforms make agents easier to deploy and manage. Shared protocols are also starting to let agents work together without a bespoke connection for every handoff. Shopify, Stripe and Anthropic each launched tools in 2026 that help agents connect to other services. [1, 2, 3]

Agents could make operations more efficient and open new revenue opportunities, but they can be costly to run and difficult to integrate, so the return needs to be demonstrated.

The direction of travel is more adoption, as enterprises look for an edge. Taken far enough, that becomes an agentic economy, with agents transacting around the clock. The pace depends not only on capability and running cost but also on public willingness to let machines act on people's behalf. The immediate test is where an agent can act safely, and who answers when it gets something wrong.

Select market signals

- Shopify announces browser AI agents can now complete purchases on its merchants' sites.[1]
- Anthropic launches blueprints to build commerce agents on Claude that enable an engineering team to get the agents running in days.[3]
- During 2026, large hyperscalers, such as OpenAI and Amazon, launch AI Forward Deployed Engineer firms to support firms in embedding agentic workflows.[4,5]
- A 2026 roadmap for the Model Context Protocol (MCP) aims for the transition into production-ready, enterprise-grade AI agent connectivity.[6]
- Stripe announces expansion of its agentic payments to include using Mastercard, Visa, Klarna and Affirm payment methods.[2]

Associated tech signals

- Agentic AI
- AI agents
- Explainable AI
- AI tokenomics
- AI regulation
- AI ethics

Trend 2

From chatbot to assistant

AI assistants can now answer questions and carry out tasks across applications, including planning the work involved. Adoption has been remarkably fast: by May 2026, ChatGPT had become the fastest app to reach around one billion monthly active users, in little more than three years.[7] What began as a chatbot can now use context to anticipate a user's needs and act on their behalf.

They are also moving beyond the screen. Meta is bringing its Muse personal agent to its latest glasses, where it can respond to what the wearer sees and act on it.[8] Across the market, new camera, audio and display models are turning eyewear into a more practical home for AI. Assistants are appearing in earbuds, watches, cars and homes too. The device matters too: an assistant in a pair of glasses encounters a different part of someone's life from one on a laptop. This is being made possible by models that can work across text, voice and images, remember earlier conversations and act rather than simply offer advice. Smaller models running on the device can do some of this without sending everything to the cloud.

How these assistants behave matters just as much as what they can do. Products such as OpenAI's Dots and Google's Gemini Spark are designed to work across apps and anticipate what might be useful, whether that is preparing a briefing before a meeting or flagging a bill.[9, 10] That could be genuinely helpful. It also brings a different set of concerns. An assistant that can see and hear what its user does may record a conversation or identify someone who never chose to be part of the interaction.

For consumer-facing organisations, this creates a practical question: what is your role when a customer turns to their assistant before they turn to you?

Select market signals

- OpenAI launches Dots, always-on personal AI agents that independently pursue goals and complete multi-step work across applications.[9]
- Meta launched its Muse personal AI agent, with around 2.8 million downloads reported in the first two weeks.[11]
- Microsoft launches Scout, an always-on personal AI agent built into the Windows ecosystem designed to function as an autonomous layer.[12]
- At Google I/O, the company made a raft of announcements that transitions its ecosystem towards proactive, agentic experiences including Gemini Spark, a 24/7 personal AI agent.[10]
- OpenClaw, an open-source AI agent that executes real-world tasks across software systems, gains rapid popularity. Its founder joined OpenAI, while the project became an independent foundation.[13]

Associated tech signals

- Agentic AI
- Edge AI
- AI chatbots and assistants
- Multimodal AI
- AI regulation
- Smart devices

Trend 3

AI's power problem

AI workloads need substantial electricity. Processing each query is becoming more efficient, but more complex and autonomous tasks use more tokens, and overall use is rising. Indeed, the International Energy Agency estimates that electricity consumption from data centres will double from 2025 levels by 2030.[14]

To meet energy needs, AI hyperscalers are signing long-term energy deals to secure supply and small modular reactors near data centres are being discussed, although regulation and community concerns remain obstacles. Nuclear fusion and solar-powered orbital data centres are longer-term possibilities being actively explored.

Some of the answer may be less power-hungry chips and smarter use of the grid, including virtual power plants that can draw on a diverse range of energy sources, such as electric vehicle batteries. Other solutions to reduce energy needs, such as biocomputing, are much further from everyday use.[15]

For most organisations, the practical question is whether energy will determine where and how quickly AI can be used, ahead of constraints such as chips or talent.

Select market signals

- OpenAI's new Jalapeño chip increases energy efficiency by delivering 1.5 to 1.9 times more AI inference work per watt than comparison systems.[16]
- The World Economic Forum reports AI developers are increasingly pursuing dedicated on-site or adjacent power supplies for AI data centres.[17]
- Annual funding in nuclear fusion climbed 69% to $4.5bn in the year to July 2026 reflecting growing confidence in the energy source.[18]
- The International Energy Agency's 2026 global energy review highlighted electricity demand rising at twice that for energy overall, partly driven by electric vehicles and data centres.[19]
- Australian Biotech firm, Cortical Labs uses 200,000 living brain cells to play the classic 1990s computer game, Doom.[20]

Associated tech signals

- Smart grids
- Fusion power
- Advanced energy storage solutions
- Alternative powertrains
- Biocomputing
- Small modular reactors

Trend 4

Trust under pressure

As AI takes on more decisions, trust depends on systems working reliably and being used safely. People also need enough information to understand what those systems do. Two risks matter most here: attackers using AI, and agents doing things their operators did not approve.

AI is already being used in fraud and financial crime. Now it may also let actors without specialist skills carry out more sophisticated attacks, while cyber operations themselves become more autonomous.

Beyond deliberate cyber-attacks, the second stand out theme is the growing issue of misaligned AI agent behaviour, where an agent acts outside the operator's instructions or policy, including actions that were never approved. During 2026 there have been cases of AI agents ‘going rogue’, often during evaluation, with some cases being widely reported.

Innovation is helping to combat these challenges and support trust in the AI era, with frontier AI models being at the forefront of defensive capabilities. Other tech such as quantum computing, privacy enhancing technologies, such as zero knowledge proofs, and digital identity may increasingly be used to support a more secure and resilient digital environment.

When both attack and defence run at machine speed, where does the human sit and what decisions need to be made?

Select market signals

- High profile incidents of misaligned behaviour by AI agents from large US AI labs.[21,22]
- President Trump and US AI leaders sign voluntary agreement to establish voluntary controls on frontier models.[23]
- Ant International, Mastercard and Visa begin development of a Know-Your-Agent interoperability framework.[24]
- Google Threat Intelligence reports a suspected AI-generated zero-day exploit and neutralises the threat before deployment.[25]
- Policymakers move to address Anthropic’s Claude Mythos 5 cyber capabilities, and the model is temporarily withdrawn whilst the Trump administration bans foreign national access.[26]

Associated tech signals

- AI cybercrime
- Deepfakes
- Digital identity
- Know your agent
- Post-quantum encryption
- Zero trust architecture

Trend 5

Robots get a brain

Humanoid robots are already appearing in workplaces, although most deployments are still pilots. Agibot's G2 is already on a live production line, while Leju Robotics and Dongfang Precision report capacity to build 10,000 humanoids a year.[27, 28]

Beyond humanoid robots, autonomous vehicles are an everyday sight in parts of China and the US, driverless trucks run commercially, and London has autonomous rides, with a licensed driver on board for now.[29]

Viral videos of athletic and combative robots have focused attention on the hardware but the more important progress is in the software, in what has come to be called physical AI. Developers are looking past large language models to world models, which learn how space, time and physics behave by simulating an environment, and many think these are the missing piece for a general-purpose robot.

If humanoid robots become cheap enough to deploy widely, questions about insurance, regulation and work become much more immediate. Who will be responsible if a robot causes harm? And how will jobs change as robots take on more tasks?

Select market signals

- Uber and Wayve launch autonomous rides in London with a TfL-licensed driver on board (September 2026).[29]
- Lidl deploys cab-less autonomous trucks for regular store deliveries in Germany.[30]
- Skild AI says its S1 robotics model enables robots to learn and perform complex tasks from a single video prompt.[31]
- The US Government bans new Chinese humanoid and dog robot imports citing national security concerns.[32]
- Japan airlines pilots humanoid robots as ground handlers at Tokyo's Haneda airport aimed at easing employee workloads.[33]
- Leju Robotics and Dongfang Precision launched an automated production line with a reported capacity of 10,000 humanoid robots a year.[28]

Associated tech signals

- Physical AI
- World models
- Computer and machine vision
- Humanoid robots
- Advanced robotics
- Autonomous mobility and logistics

Trend 6

Quantum sets a deadline

Quantum computers approach problems very differently from conventional machines. Useful systems are still being developed, but organisations already have work to do on encryption.

The past few years have brought progress towards stable, useful quantum computers with companies and governments investing heavily. Some expect the technology to prove as transformative as AI. The immediate work is to make quantum systems bigger and more reliable, then prove where they can be useful.

The application of quantum computing is expected to be wide-ranging, excelling at optimisation and problems involving very large datasets and dimensions. Examples of potential use cases include more effective drug discovery, optimisation of logistics and supply chain problems, optimisation of investment portfolios, and a step-up of AI model training and capability.

The security question is more immediate. IBM, Microsoft and Google have all put dates before the end of the decade to move to post-quantum cryptography. [34,35,36]

That does not mean today's encryption fails by 2030 but it does mean migration takes planning and should not wait for a breakthrough.

Select market signals

- The US Government signs Executive Orders to accelerate domestic quantum computing development.[37]
- IBM announces it will invest more than $10bn into quantum computing over the next 5 years, supporting the roadmap to a large-scale fault-tolerant quantum computer by 2029.[38]
- Microsoft unveils their 2nd generation Majorana quantum chip which they claim is 1,000x more stable and could accelerate the timeline to a scalable, fault-tolerant quantum computer to 2029.[39]
- Google targets 2029 to transition to post-quantum cryptography citing progress in quantum computing.[36]
- The UK Govt announces a £2bn investment package to build large-scale quantum computing infrastructure in the UK by the early 2030s.[40]

Associated tech signals

- Quantum computing
- Quantum communication
- Cloud computing
- Quantum encryption
- Digital twin tech
- Post-quantum encryption

Trend 7

Healthspan for the ages

New technologies could help people live longer in good health, extending what researchers call healthspan.

AI helps researchers make sense of large datasets, supporting earlier risk detection and speeding up drug discovery. AI-powered systems are increasingly supporting personalised medicine, enabling treatments to be tailored to an individual's genetics, lifestyle and health history. Quantum computing is set to extend these types of capability.

Surgeons are using AI to help with complex operations, but this is still some way to go until this becomes routine. Robotics may also eventually help patients recover mobility or give older people more support, although these applications remain at an earlier stage.

Wearables already offer basic health features but with more capable AI running on the devices, they could offer more personalised advice and help people manage their health. Services such as ChatGPT Health show the early shape of a more proactive healthcare assistant.

Together, these technologies could detect disease earlier and help people stay well for longer but they also raise questions that reach well beyond medicine. If people live longer in good health, what happens to working lives, retirement, insurance and the products built around a shorter old age? And who do people trust with the most personal data they have?

Select market signals

- In clinical tests, Microsoft's MAI-DxO system achieves higher diagnostic accuracy than generalist doctors for complex medical cases.[41]
- In a world first, surgeons removed a brain tumour with assistance from AI that analysed hidden vessels and nerves in real time.[42]
- Google unveils SensorFM, a foundation model for wearable health, trained on a trillion minutes of sensor data from 5m people.[43]
- AWS launches Amazon Bio Discovery, an AI-powered application to speed up drug development that can be used without coding knowledge.[44]
- Healthcare plugins for ChatGPT, Claude and Amazon aim to offer more effective health support to users (but do not replace health professionals). [45, 46, 47]

Associated tech signals

- AI diagnostics
- AI drug discovery
- Robotics in surgery
- Wearable health tech
- Domain-specific language models
- Personalised medicine and genomics

Trend 8

Orbit as infrastructure

Commercial space activity has grown, helped by reusable launches and large satellite networks. One catalogue counted more than 18,000 payloads in orbit in October 2026, including inactive ones.[48]

Beyond exploration, companies are building services around launches, satellites, orbital infrastructure and the data they produce. If those services become commercially viable at scale, with more satellites in low Earth orbit (LEO), this could lead to improvements in connectivity, environmental monitoring and support research.

Microgravity may make it possible to manufacture higher-quality semiconductors, fibre optics and pharmaceuticals in orbit. There is an expectation that it could also lead to new advanced materials and alloys that are simply not possible to develop on Earth.

For most organisations, the near-term question is practical. As critical services come to depend on satellite links, is that dependence understood, and is it resilient? The longer-term one is who gets to own the infrastructure of orbit.

Select market signals

- Crunchbase data shows that in Jan-Aug 2026, there was a record $20.3bn of global seed to growth stage funding into the space and satellite-related sectors.[49]
- Skyroot Aerospace successfully launches Vikram-1, India's first privately developed orbital launch, placing customer payloads into low Earth orbit.[50]
- Amazon agrees to acquire Globalstar for $11.57bn to expand its satellite network.[51]
- The UK Goverment allocates £500m to back high-growth space tech to scale and compete on the international stage.[52]
- Axiom Space successfully launches the first two orbital data centre nodes to lay the foundation for space-based cloud computing.[53]

Associated tech signals

- Space tech
- 6G connectivity
- Nanotechnology
- Metamaterials
- Advanced robotics
- Autonomous mobility and logistics



02.

Double-clicks

AI agents are moving from answering questions to taking action. The next three sections look at what agents cost, how they might change shopping and why people may hesitate to trust them.

1

The economics of autonomy



2

Agentic commerce takes shape



3

Inside the AI trust paradox



The economics of autonomy

The context

Falling token prices made early AI pilots look inexpensive, but the total bill looked different once agents began making repeated calls.

Token prices are continuing to fall, with Gartner predicting inference costs for a large AI model to fall by more than 90% between 2025 and 2030.[54] But as models have become more capable and agents turned an individual task into repeated model calls and tool use, often with retries along the way, the bills have gone up.

Employees learned that burning tokens looked like productivity, and organisations had to crack down on it. Uber reportedly capped AI coding spend per person after getting through a year's budget in just four months.[55] The cheaper the unit, the bigger the bill: an old pattern in a new setting. Scrutiny then followed. BCG found nine in ten chief executives saw some value from AI but only 14% defining P&L impact for all AI initiatives.[56] KPMG found 42% of UK leaders could not fully see what they were spending.[57]

What comes next is less dramatic and more useful: the discipline of buying AI well.

Recent developments

- IBM announces Apptio AI Value & ROI, a set of capabilities that connects AI spending, including token costs, to measurable business outcomes.[58]
- Revefi launches its AI FinOps, Observability Token Economics suite allowing real-time spend protection and end-to-end AI attribution.[59]
- Nvidia launches NeMo Switchyard, an open-source AI model router for which LangChain achieved a 74% cost reduction in multi-turn agent tasks.[60]
- A BCG Executive Perspectives survey finds that 60% of companies are not achieving material value from AI despite significant investment.[61]
- Uber reportedly introduced a $1,500 monthly spending cap per employee for each AI coding platform after reports that it consumed its annual AI budget within the first four months of 2026.[55]
- In early 2026, ‘tokenmaxxing’ – where employees maximise use of AI tokens to signal productivity and performance – emerges, followed by a crack down on the practice.[62, 63]

The economics of autonomy

An exploration

Whilst a chatbot can answer a question with one model call, an AI agent may plan, use tools, check and retry, all for a single task, resulting in many calls to a model. Falling token prices made using AI look cheap, so usage ran ahead of budgets.

Incentives then made it worse, where token use became a proxy for effort, people maximised it. And when nobody owned the bill, nobody had much reason to challenge it.

Yet, the wider costs were always there: cloud and storage underneath the tokens and, outside the technology budget, the data clean-up, integration, controls, governance, training and redesign of the work itself.

Cloud computing went through a similar cycle a decade ago: metered services looked cheap, then the bills ballooned. FinOps emerged to make consumption visible and assign someone responsibility for it. AI is repeating that cycle, only faster, and with a unit of cost that enterprises are still learning to understand.

As agents take on more work, the choice of model starts to matter to the overall bill. A summary might work well on a smaller model but approving a



The economics of autonomy

An exploration



payment is a different proposition. Smaller and open-weight models could also handle some narrower tasks, including work run on an organisation's own infrastructure or on a device.

The harder question is where the line sits. A cheaper model may save money on one task and create more checking or correction on another.

How much of an agent's work really needs the most capable model? And how would an organisation tell when a cheaper one is good enough? But the price of a model call only tells part of the story.

What does it cost to finish the job, including any checking and corrections? Caching, batching and shorter context might bring that cost down, though the benefit will depend on the work.

There is also the question of who can see the total bill as agents spread across teams. Some benefits, such as lower handling costs, may be fairly easy to count, but faster delivery or a loss avoided is harder to put a figure on. That leaves enterprises asking what evidence would show whether an agent is worth its cost compared with the way the work is done today.

The economics of autonomy

Reflections to consider

Who owns the bill?

Someone should own a complete view of AI running costs, including what drives them and whether each team's usage pays its way. Without an owner, usage can grow long before anyone can explain the bill.

What does a completed task cost?

Price per token flatters cheap attempts. Cost per task done to the quality you need, including the retries and the human checking, is the number that bears comparison with the process it replaces.

Which model earns its place?

Use a frontier model where its extra capability earns the higher cost. Open-weight and small models can do bounded work for a fraction of the cost, if there are evaluations to prove the output is good enough.

What happens when the price collapses?

If the price per token falls by 90% and usage rises a hundredfold, spending rises tenfold in this illustration. Would the extra capacity be worth the cost?

Agentic commerce takes shape

The context

Online shopping has made it easier to find products and pay for them, but people still usually make the final choice. AI agents could take on more of that work; some services can already finish a purchase once the buyer confirms it.

A further step would be asking an agent to find something within a budget, choose the seller and pay on your behalf. The buying part is agentic shopping; the transaction is agentic payments. Both sit within agentic commerce.

Sellers would then face a different problem. Their offer may have to get past an agent before a customer even sees it, and the customer may not know why that offer was chosen. Anthropic's Project Deal shows why the quality of the agent could matter too: in a small experiment with 69 staff, stronger agents got better prices, while people using weaker ones were unaware that they got worse prices.[64] We cannot assume everyday shopping will work the same way. But the agent someone uses might affect the deal they get.

Recent developments

- Meta and Sierra develop Personal Agent Protocol with partners, an open standard to define a secure, consistent way for personal AI agents to interact with businesses on customers' behalf.[65]
- TikTok rolls out an AI shopping assistant and one-click checkout, allows users to conduct the full shopping journey within the TikTok platform.[66]
- Amazon blocks Meta's personal AI agent, Muse, from accessing its platform, citing Conditions of Use violations, similar to previous legal struggles against similar agents.[67]
- The American Arbitration Association launches the Legal Context Protocol, an open standard to make legal terms, consent, and dispute rules discoverable and verifiable for AI agent transactions.[68]
- Google introduces Universal Cart, an intelligent shopping cart that acts as an agentic hub for shopping across the Google ecosystem and simplifies the checkout process.[69]
- Anthropic's Project Deal experiment demonstrates the ability of AI agents to operate autonomously on behalf of buyers and sellers in a marketplace.[64]

Agentic commerce takes shape

An exploration

Online shopping and smartphones made it easier for buyers and sellers to find each other. AI could change what happens after they connect: an agent may recommend a purchase and arrange payment, sometimes completing the transaction with little human involvement.

For now, agents are most visible in product discovery and checkout although some services can complete a purchase once the buyer confirms it.

Agentic shopping may spread as the technology improves and shared protocols develop, although its reach will also depend on whether the rules around it earn trust. And the sift could happen quickly: AI is becoming already a primary channel for people and organisations alike.

For consumers, that could mean smoother, more personal experiences and the option of delegating routine purchases, such as groceries, altogether. For businesses, agents could extend reach and keep commerce running around the clock as well as reducing the work needed to complete a sale. It also means a new intermediary between businesses and customers.

Adoption is uncertain, however. Some people may not trust an agent to shop for them; a question explored in the AI trust paradox double-click. There



Agentic commerce takes shape

An exploration



will always be purchases people want to be involved in, so agentic approaches are more likely to complement existing shopping approaches rather than entirely replacing them.

Sellers might reach more people through agents but have less direct contact with customers. AI aggregators may impose their own costs and standards, much as app stores did.

Agentic shopping may also demand a change in strategy simply to be visible: to be found by AI aggregators and included in what they show. Beyond shopping, agentic AI reaches into the wider world of commerce, where organisations transact through agent channels that operate continuously and at machine speed. NatWest, in conjunction with five other institutions has published high-level customer centric principles to support trust and scale in agentic commerce.[70]

As the world works out how agents fit alongside employees, the consequences for commerce and society could be significant.

The timing remains uncertain, but businesses can already explore where an agent might sit between them and a customer or supplier.

Agentic commerce takes shape

Reflections to consider

Operational readiness

Taking part in agentic commerce is likely to need changes in data, systems and governance, so that an organisation's own workflows can talk to the agents operating across these ecosystems. How far is the current estate from being able to do that?

How agents see brand

If an agent chooses what a customer sees, how does a business make sure its products are considered? A customer might never reach the brand's own interface. Should the way we measure brand success change?

Choosing agent channels

AI chatbots and other AI interfaces are new channels, and they diversify where customers are. Which deserve priority, and what does good presence and engagement look like in each?

New markets for agentic commerce

The internet opened new markets, products, services and capabilities. Agentic commerce may do the same. Where are the new areas of business, and what would it take to be early to them?

Inside the AI trust paradox

The context

Using AI is one thing; letting it decide for you is another. In an EY survey, 74% of UK adults had used AI in the previous six months, but only 14% were comfortable with a fully autonomous agent.[71] That gap is worth understanding before organisations ask customers to delegate decisions.

People may worry about errors, privacy or jobs but these figures do not tell us which concern explains the gap. Reports of agents taking unauthorised actions raise further questions about safety.[21,22]

Customers may welcome AI's convenience while hesitating to let it make decisions for them. Organisations need to understand those limits and show that their systems are reliable and well governed.

Recent developments

- High profile incidents of misaligned behaviour by AI agents from large US AI labs.[21,22]
- Spotify introduces AI persona badges to identify artist profiles that represent a non-human persona, with music by AI personas excluded by default from recommendations.[72]
- Menlo Ventures' 2026 State of Consumer AI found that accuracy (45%), trustworthiness (40%) and privacy (36%) were the most important factors for consumers choosing an AI tool.[73]
- In the Usercentrics State of Digital Trust 2026 report, 52% of respondents said they would pay an average 7% more for AI transparency. The report also found that 47% had acted on AI trust concerns in ways that affected companies' revenue.[74]
- A 2026 Ernst & Young survey finds that, amid increasing AI usage, trust in institutions governing of AI and managing AI-related data remained low, with only 14% of UK respondents comfortable relying on fully autonomous AI agents.[71]

Inside the AI trust paradox

An exploration

AI use is growing, but surveys suggest trust is weakening. Gallup's survey of 14- to 29-year-olds in the US found weekly use steady at around half, while excitement fell and anger rose.[75] The reasons are uncertain; concerns often cited include fewer opportunities in entry-level jobs, loss of skills, feeling pushed to use AI, limited transparency and its environmental impact.

Reports in 2026 of AI agents taking unauthorised actions, including access to third-party systems, prompted wider debate about safety and existential risk.[21,26] They also raised a practical question: can controls keep up with model capability?

Some people are trying to spend less time online. That may reflect screen fatigue, but we do not have evidence that distrust of AI is driving it.[76]

For businesses and policymakers, the practical question is how willing people will be to use AI services as trust changes. The answer is hard to predict.

How trust develops may depend on how quickly AI improves and whether it becomes reliable enough for decisions people care about.



Inside the AI trust paradox

An exploration



In one future, development accelerates towards artificial general intelligence, systems with human-level capability or better. Reliability and usefulness might lift trust but the disruption to society, mass job displacement for instance, could just as easily undermine it, as today's fears turn into experience.

In an alternative future, progress is slower and capability more modest. Reliability remains a drag on trust, but society keeps pace, deployment is more balanced and people acclimatise. Trust could end up higher in the slower world than in the faster one.

Neither path is a forecast.

Organisations can test where AI works reliably for their customers and how to govern it. The useful work now is to test services with customers rather than assume a broad survey speaks for them.

Inside the AI trust paradox

Reflections to consider

Transparency in how AI is used

Where the models allow, openness about where AI is used and the effect of its decisions can help customers understand it. Clear explanations may help customers judge a service, but they will not make up for decisions that do not work for them. How transparent is the organisation today, and how much more could it be?

Providing agency to customers

Customers may be more willing to use AI if they can choose how much it is involved and what data it uses, for example in a personalised service. Those choices would let people with different levels of trust use the service differently.

Governance and ethics

Trust rests on AI making decisions that work for customers and treat them fairly. That depends on the strength of oversight over AI workflows and on how AI is used. Is governance keeping pace with deployment?

Standards for reliability and performance

Clear benchmarks for AI outputs support outcomes customers can trust. Those standards may need to be higher than for people, because customers rarely forgive a machine's mistake the way they forgive a person. What is the bar, and who sets it?



03.

Our methodology

How we identify the eight tech-led trends.

We start by collecting signals from emerging technologies and the policy, markets and behaviour around them. This research covers where a technology might have a real-world effect.

The NatWest Open Innovation team draws on its research into emerging technology and its work with innovators. Colleagues across the bank add perspectives from their markets and from customers.

Finally, we check our list against a range of comprehensive third-party technology trend reports. That comparison helps us understand how our conclusions relate to the wider consensus.

We score each potential emerging tech, and the developments around it, on five attributes:

1. *Novelty:* the level of change expected from the area of development, relative to what exists today.

2. *Impact:* the extent to which it is expected to yield real-world utility or consequences.

3. *Depth:* how much research and practical development sits behind it.

4. *Momentum*: the speed and direction of progress.

5. *Timeliness*: the time horizon over which it is expected to be realised or adopted at scale.

We score each signal and put it on the radar. We then look for connections across the signals, drawing on what colleagues and external sources tell us. Those connections inform the eight trends.

The market signals on each trend page are from 2026 unless another date is given. The market signals provide examples of each trend. The trend itself draws on the wider body of evidence. Regulatory decisions provide different evidence from research or company demonstrations, so we indicate which is which where it matters.

We use a foresight approach to inform our double-click insights

Each deeper discussion uses four steps to explore how a trend could develop

1. Horizon scanning

As we are exploring tech-led disruption, the starting point is identifying tech signals. We draw on a variety of sources: the disruption radar, market signals, internal and external surveys and data relevant to the double-click, and the views of experts from across the bank.

We also identify the non-tech signals, government policy or geopolitics for example, that could affect the theme, so that the view of how change might unfold is a rounded one.

2. What could shape the outcome

We identify the trends with the greatest influence on how the topic might unfold over time: (a) the tech trends identified in this report, (b) tech and non-tech trends happening today that might change the theme, (c) visions of the future associated with it, and (d) potential barriers to change.

We consider when an influence may take effect and how much it could change the outcome, alongside the uncertainty around it.

3. Scenario builds

We sketch several plausible ways the issue might develop and note what would have to happen in each case. We also consider unlikely events that could change the outcome materially.

4. Strategic insights

We compare the scenarios to find recurring outcomes and the decisions that remain sensible under uncertainty. We consider:

- Notable commonalities across scenarios.
- The most significant uncertainties, and the factors that most influence them.
- Potential changes to the competitive landscape and to customer and user needs.
- Opportunities and threats associated with the scenarios.
- Strengths and weaknesses that might help or hinder taking the opportunities or addressing the risks.
- Strategic options that would work across scenarios, and those that would need to flex.

Sources and references

1. Shopify
2. Stripe
3. Anthropic
4. AWS
5. OpenAI
6. MCP
7. Reuters
8. Meta
9. OpenAI
10. Google
11. TechCrunch
12. Microsoft
13. Reuters
14. IEA
15. Final Spark
16. OpenAI
17. WEF
18. FIA
19. IEA
20. ACS
21. Reuters
22. Anthropic
23. Forbes
24. Ant International
25. Google
26. Reuters
27. Agibot
28. China Biz Insider
29. Uber
30. Reuters
31. Skild
32. BBC News
33. JAL
34. IBM
35. Microsoft
36. Google
37. White House
38. IBM
39. Microsoft
40. GOV UK
41. AI News
42. UCL
43. Google Research
44. AWS
45. Anthropic
46. Amazon
47. OpenAI
48. Orbital Radar
49. Crunchbase
50. ISRO
51. Amazon
52. UK Space Agency
53. Axiom Space
54. Gartner
55. TechCrunch
56. BCG
57. KPMG
58. IBM
59. Revefi
60. NVIDIA
61. BCG
62. Gizmodo
63. Financial Times
64. Anthropic
65. Sierra
66. TechCrunch
67. Bloomberg
68. AAA
69. Google
70. NatWest
71. EY
72. Spotify
73. Menlo Ventures
74. Usercentrics
75. Gallup
76. The Guardian

Disruption Watch Trends H1 2027

For further information about this report, please contact:

OpenInnovation@natwest.com



We've prepared this report to share our views and insights on emerging tech-led trends. The trends and views we discuss may change as the technology landscape evolves. It is not investment, legal, tax or other professional advice and it should not be used as the basis for any business or investment decisions. NatWest Group accepts no responsibility for any loss arising from its use. AI-generated images used for illustrative purposes.

Copyright @ 2026 NatWest Group