Press contact:
Elsa Estager Bergerou
Tel.: +33 6 59 62 55 13
E-mail: elsa.estager-bergerou@capgemini.com

TechnoVision: Top 5 Tech Trends to Watch in 2027

Paris, November 17, 2026 - Capgemini unveiled today its *TechnoVision Top 5 Tech Trends to Watch in 2027*, highlighting the technologies that are expected to reach an inflection point in the next year. As companies aim for greater autonomy across their processes and operations, AI is entering a new frontier, moving beyond its role as an 'assistant' towards systems that can now act and execute. In 2027, this shift will reshape enterprise operations, cybersecurity, R&D and commerce, opening new possibilities while raising challenges around governance, safety and resilience.

"*Last year, Capgemini's Top 5 Tech Trends highlighted how AI was becoming embedded across enterprise architecture, software engineering and business operations*," explains Pascal Brier, Chief Innovation Officer at Capgemini and Member of the Group Executive Committee. "*With AI agents acting autonomously, AI is reaching a turning point. As we look ahead to 2027, the main challenge for organizations will be how to scale autonomy without losing control, which will require managing interactions between agents, containing unexpected outcomes and responding to threats at machine speed. Beyond these challenges, agentic AI will reshape established ways of working, from accelerated innovation in R&D to reinventing interactions between businesses and consumers.*"

Technologies to watch in 2027:

1) **The agentic enterprise - multi-agent systems: Agents interact in ways nobody designed**

Today, most AI agents and agentic systems are simple enough to be tested before they go live. But as their number grows, it becomes impossible to test every interaction between them. For example, a company running 15 AI agents has 105 possible pairs; with 1,000 agents and 200 agentic systems, this rises to almost 720,000. At that density, agents could produce behaviors that nobody designed. Each agent can make a reasonable decision based on its own objective and available information, yet those decisions can interact, trigger reactions from other agents and ultimately lead to an unintended outcome at system level. Every agent can be right, while the system is wrong, and the business consequences can be serious.

Why it matters: Future control will need to operate at three levels: a harness around each agent, a control plane for each agentic system, and an enterprise control plane across all of them. In 2027, enterprise control planes will need to incorporate a register of every agent and agentic system, an identity and permissions for each one, monitoring of outcomes delivered and consumption, visibility into how agents interact, limits on what they can do together, and a way to stop them.

2) **Trusted autonomy & safe AI: From explainability to containment**

Ensuring safe AI has traditionally relied on recording how AI reaches decisions, limiting the access of AI agents and agentic systems, and checking that every decision remains within business rules. But this approach is coming under pressure as some new models no longer show all their reasoning steps. As AI systems become more sophisticated and operate at greater scale, monitoring all their thinking becomes more difficult. Safe AI therefore means accepting that errors and incidents can happen and deciding how far each one can spread. In practice, this means that organizations need to watch for warning signs, such as an agent using an unauthorized tool or another one losing part of its instructions. It could also be a sudden spike in cost or response time, or a system going down. When these signs appear, the agent must be checked or taken down.

Why it matters: In 2027, organizations will need to focus on AI containment instead of explainability. They will also need to set up enterprise control planes covering all agents and agentic systems, including visibility into agent-to-agent interactions and early signs of technical errors.

3) Machine-speed cyber: Cyber attacks now move at machine speed

An AI agent can now find software flaws on its own and turn them into real attacks, with no human involved, even if they were not designed to do so. On the attack side, the risk is growing as AI increases the capabilities available to identify vulnerabilities and build exploits, including attacks capable of taking control of critical systems. On the defense side, the attack surface keeps expanding: every AI assistant, agent and agentic system deployed creates a new potential entry point, while intrusion detection is currently not improving at the same pace. As attacks become increasingly automated and industrialized, a cyber defense that waits for a human analyst to react has already lost.

Why it matters: In 2027, as cyber attacks increasingly move at machine speed, security operations will need to do the same. Security operations centers will need machines to triage and respond, while humans set the rules, deploy agentic systems and supervise. This also means moving toward a new generation of security operations centers, enabling detection within minutes of a new threat and much faster remediation.

4) Agentic R&D: Agentic R&D breaks the capacity limit

In R&D, the capacity to research, screen and validate options has traditionally been limited by available resources and testing capacity. AI changes that equation across two key phases: exploration, finding the right idea for a product, and execution, turning that idea into a product. In exploration, organizations have historically been limited in how many options they could test. AI removes that ceiling, enabling more possibilities to be explored across molecular products such as drugs, chemicals and materials, as well as physical products such as vehicles, machines and consumer goods. In execution, AI is entering every step, with agents writing, testing and fixing code, and AI increasingly supporting design, simulation, testing and product lifecycle activities.

Why it matters: In 2027, AI will enable organizations to explore and evaluate a much wider range of options while accelerating execution. In exploration, this means more capabilities to faster identify and better assess promising possibilities such as molecule discovery. In execution, engineers will increasingly be AI-assisted, reducing costs and accelerating time-to-market. The R&D operating model will need to evolve around AI, agents, data and AI-native tools.

5) Agentic commerce: Agents on both sides of the transaction

AI agents are changing both sides of commerce: buying and selling. On the buyer side, consumers can ask an AI assistant to find products based on criteria such as price, availability and delivery, with the agent checking multiple retailers in seconds for product, price or delivery details. On the seller side, retailers are beginning to use agentic systems to run their operations, enrich catalogues, adjust prices and promotions, refill stock, and handle customer service and returns. Agents are rapidly becoming part of the transaction itself.

Why it matters: In 2027, organizations will increasingly need to make their businesses readable by AI. This means providing clean data on products, prices, stock and delivery, as well as a way to authenticate the agents they let in and keep bad agents out. On the seller side, agentic systems can bring greater cost efficiency and speed to commerce operations. At the same time, as AI shapes how consumers discover, evaluate and purchase products, brands will need to both ensure their products remain visible and relevant to AI assistants while also continuing to build customer preference, loyalty and trust.

About Capgemini

Capgemini is the business transformation partner for enterprises in the age of AI. We help organizations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.

Make it real | www.capgemini.com