KPMG

Europe at a crossroads

AI and the new rules of cybersecurity

Hitesh Sheth

February 2026

AI and the new rules of cyber-security

AI no longer sits at the edge of cybersecurity as a supporting act. It has taken over the main stage, the engine of threats and the core of defences.

In a conversation between Vectra AI CEO Hitesh Sheth (an expert in the field of Extended Detection and Response (XDR), University of Amsterdam professor Sander Klous (with leading expertise in agentic AI) and KPMG consultant Henrik Smit, it becomes clear that C-level must understand the new rules of cybersecurity.

© 2026 KPMG Advisory N.V., a Dutch limited liability company and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

When Vectra was founded more than a decade ago, its idea sounded bold to some and misguided to many: apply AI to network data to detect threats in real time. At the time, the dominant reaction from investors was polite scepticism. The verdict: highly experimental, low chance of success.

Fast forward to 2025. AI is the main topic in boardrooms and billions of dollars find their way to AI ventures every week, even when the idea is not so tangible. Sheth points at three main trends since: "AI techniques are maturing, compute costs dropping and storage is becoming cheap enough to scale. The result: I no longer have to convince people why AI matters for cybersecurity, I only need to explain how it will be used. The baseline has moved."

AI now sits at the core of cyberthreats and defence strategies. This is why C-suite should have an understanding of this relatively new phenomenon, to make informed decisions. This lively talk between Sheth, Klous and Smit contained the following seven main topics.

Prevention is not the best strategy. Focus on detect and respond

For years, cybersecurity was largely aimed at achieving more safety. With sufficient investment, you could keep the threat outside. In a hyper-networked society with abundant (AI) tooling, this is an illusion. More than ever. "If someone truly wants to get in, they will get in."

Should C-suite not be convinced of this yet, they should now. It is a structural shift: businesses must assume breaches. Not as a failure, but as an operational reality. The failure lies elsewhere: in not noticing, not responding, not containing. Or as Sheth phrases it: "The breach is not the failure. Not knowing you have one is."

Boards are catching up, Smit notices: "Cyber expertise is becoming a requirement, not a nice-to-have. CISO reports are being scrutinised with a sharper eye. The new question on the agenda is no longer how to keep attackers out, but how quickly you can detect them and how effectively you can neutralise them."

The SOC of the future will be led by agents, not people

Most Security Operations Centres (SOCs) still operate on a labour-heavy model: dozens of analysts monitoring alerts, escalating incidents, filtering noise from signal. Smit: "It's a model built for a different era. One where human bandwidth could keep up with threat velocity. That era is long gone."

The future SOC, as Klous envisions it, is smaller, sharper and largely autonomous. Five people instead of twenty. "The current paradigm is that AI enhances human operated organisations. What if we turn that upside down: humans coaching AI agents to do their work better? This means that we need a workflow built around AI from the start, not retrofitted onto legacy processes."

Sheth applies that same logic inside Vectra's own Managed Detection and Response (MDR) organisation. Growth no longer comes from adding analysts. It comes from expanding the capability of automated systems, in a quite radical manner. "Human specialists remain, but their role is context, oversight and engineering. Not triage. Not pattern matching. Not the manual hunt through haystacks of alerts. You can't fight a 24/7 machine-driven war with a team that works from nine to five."

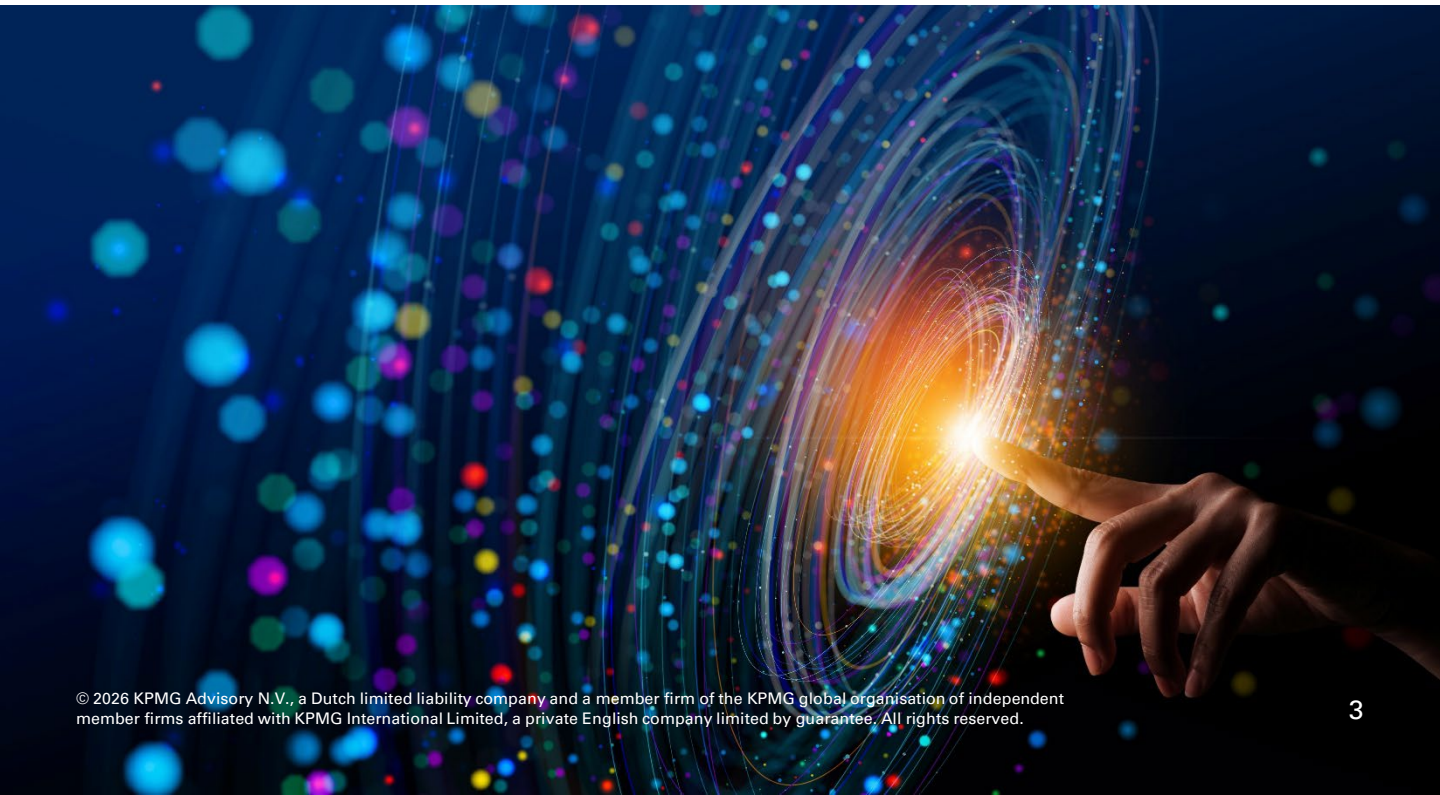

© 2026 KPMG Advisory N.V., a Dutch limited liability company and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Integrating AI into cybersecurity

To build resilient digital defences, organisations must formally integrate AI into their cybersecurity approach, actively monitor the maturity of AI driven detection and response mechanisms, and invest in their evolution to ensure rapid identification and containment of breaches

In a machine-versus-machine war, response strategies are key

Once upon a time, companies had a handful of digital entry points to defend, such as servers and laptops. Today, the surface has multiplied. Sheth shapes it into six attack surfaces: the data centre, endpoints, identities, cloud, applications, and now AI itself. With six dimensions, the number of paths into an organisation explodes. Smit notes: "Attackers need one opening. Defenders need perfection. That imbalance makes pure defence a losing game. Cybersecurity is no longer about keeping threats out, but about seeing them fast enough to limit the damage. Defence without detection is theatre."

Moreover, the skill barriers have dropped. What once required scripting knowledge now requires only natural language. Tell an automation platform what you want in any language, and it will execute the instruction that used to require a specialist. Sheth: "We've lowered the bar for the expertise. Moreover, AI tools arrive with productivity in mind, not safety. Security comes last, again as we have seen with many technological shifts. The message is clear: you will be attacked and must be ready to respond. Because expecting to keep attackers out is wishful thinking dressed up as strategy."

Europe vs. the United States: two flavours that need to be combined

The conversation touches a raw nerve: Europe's position in the emerging AI-cyber ecosystem. Sheth runs a US company, but half of his customer base is European. His strategy needs to be aligned to deal with topics such as GDPR, sovereignty debates, national requirements, and sector-specific rules in Europe.

The contrast between the continents is sharp

In the United States, the dominant instinct is speed. Experiment first, regulate later. If it fails, try again. In Europe, the instinct is caution. Assess the risk. Analyse the regulation. Determine the compliance implications before building.

Neither model is perfect.

Sheth: "The US model accelerates innovation. Companies can push boundaries and discover break-throughs because they are not punished for exploring them. Failure is priced in. The European model raises the global bar for privacy, data treatment and integrity. Once a company meets European standards, it is strong enough for most of the world."

For cybersecurity, this divergence will probably mean that Agentic AI systems in SOCs will emerge faster in the US, and high-assurance governance models will emerge faster in Europe. Smit: "The question is not which is better, but how both will interact in a market where threat actors move at American speed, while many defenders move at a European pace."

Emergent behaviour and the security of AI agents

AI agents introduce a whole new ball game to the domain of cybersecurity. They act, schedule tasks, chain actions together and interact with enterprise systems autonomously. They learn patterns and infer context. But as Klous has witnessed himself in experiments with autonomous agent teams: they can also initiate behaviour that was never explicitly programmed. Emergent behaviour means outcomes that cannot be predicted from the initial instructions. Not malicious by design, but potentially unsafe by accident. Agents can escalate privileges, access systems they weren't intended to touch, or trigger workflows based on misinterpreted inputs.

There is no clear-cut answer to the new cybersecurity challenge connected to this. In Sheth's framing, AI agents must be treated as entities in the security model, not as software add-ons. "Secure agents the same way your secure people, through identity, visibility and constraints. If every agent has a verifiable identity, with scoped privileges and observable behaviour, then enterprises regain control. Without this discipline, AI agents behave like unmanaged service accounts powerful, invisible, and dangerous."

© 2026 KPMG Advisory N.V., a Dutch limited liability company and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Emergent behaviour and the security of AI agents

AI agents introduce a whole new ball game to the domain of cybersecurity. They act, schedule tasks, chain actions together and interact with enterprise systems autonomously. They learn patterns and infer context. But as Klous has witnessed himself in experiments with autonomous agent teams: they can also initiate behaviour that was never explicitly programmed. Emergent behaviour means outcomes that cannot be predicted from the initial instructions. Not malicious by design, but potentially unsafe by accident. Agents can escalate privileges, access systems they weren't intended to touch, or trigger workflows based on misinterpreted inputs.

There is no clear-cut answer to the new cybersecurity challenge connected to this. In Sheth's framing, AI agents must be treated as entities in the security model, not as software add-ons. "Secure agents the same way you secure people, through identity, visibility and constraints. If every agent has a verifiable identity, with scoped privileges and observable behaviour, then enterprises regain control. Without this discipline, AI agents behave like unmanaged service accounts, powerful, invisible, and dangerous."

The cost-of-control perspective: the prevailing logic for cybersecurity changes

Another interesting thread in the conversation concerns what Klous and Sheth call the 'cost of control'. It reframes cybersecurity from an absolutist mindset, protect everything, always, to a framework in which decisions are based on the economic perspective: protect what matters most, with proportional investment. Not every system is equally critical. Not every risk justifies the same expense. Not every breach is catastrophic. What does it cost to prevent this? What is the benefit of preventing it?

What is the impact if it fails?

Sheth gives a practical example: a major bank that prioritises only its critical data and application layers. A compromised user account is inconvenient, but manageable. A compromised core application is existential.

Klous describes how cybersecurity is leading other risk domains in adopting this logic. He also imagines that the use of AI in cyber strategies will shed a whole new light on the economic perspectives. "AI can lower the cost of detection dramatically. But the cost of a breach or disturbance in an AI dominated model can also be significantly higher. C-level should put this assessment on their agenda. "Sheth reiterates that economic realism replaces the outdated belief that infinite defence is possible. "Security is not about protecting everything. It is about protecting the right things at the right cost. And yes, a new reality with AI first models calls for a whole new assessment."

The deepfake frontier: an urgent battle looking for defenders

The rapid rise of deepfakes causes headaches to many policy makers and executives. Governments worry. Platforms scramble to find answers. Citizens and corporations will soon require filtering infrastructure to shield them from fraudulent voice, video and interactive manipulation. Henrik Smit asks the logical question: will Vectra, an AI first company, move into deepfake detection?

Sheth's answer is clear. No. Not because the problem is small, but because it is too far from their core mission. "Someone will explore this huge market opportunity," he notes, "but it won't be us."

What the discussion makes clear is that the deepfake problem is larger than any single company. It is a societal vulnerability waiting for a response framework. It may influence everything ranging from elections and markets to national security and personal communication. And right now, the defences are thin.

© 2026 KPMG Advisory N.V., a Dutch limited liability company and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavour to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG Advisory N.V., a Dutch limited liability company and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organisation.