

Five AI developments of 2026 that change how boards work

1. Executive summary

Switzerland has no AI Act (yet), however directors are responsible for : oversight (Art. 716a CO), due care (Art. 717 CO) and personal liability (Art. 754 CO). In 2026, AI became a board question of value, control and capability. Boards which change how they work can create value and contain risk; those which do not change will oversee spending they cannot evaluate and risks they cannot see.

The logic

- 1 Money is spent faster than its proof of value materializes.** AI budgets double while only 37% of organisations report any EBIT impact (Development 1).
- 2 Autonomy and Capability are advancing faster than Control.** Agents act autonomously, frontier models find software flaws (bugs) at scale, and trillion-parameter AI models now run on a local PC (Developments 2–4).
- 3 The board's own working methods are the weakest link.** Directors already use AI, but few boards have rules, skills or information to govern it (Development 5).

#	Development	Opportunity and threat	Shift in board working method
1	AI value gap	Opportunity: the 6% of companies that redesign work capture 5%+ of EBIT from AI. Threat: spending rises on belief, not evidence.	From approving AI budgets to steering value with stage gates and a value dashboard
2	Agentic AI in operations	Opportunity: end-to-end process redesign. Threat: Accountability? Automated decisions nobody owns; EU rules deferred, not cancelled.	From policy statements to an agent register, autonomy limits and a named responsible executive (owner)
3	AI-driven cyber risk	Opportunity: defenders can find flaws first. Threat: cheap, automated vulnerability discovery for attackers.	From an annual cyber briefing to metrics, tabletop exercises and supplier-risk challenge
4	Frontier AI running on your desk	Opportunity: sovereign, in-house analysis of highly sensitive and confidential data. Threat: shadow ¹ AI with frontier power.	From buying tools and services to approving infrastructure principles by data class
5	Directors' own AI use	Opportunity: faster preparation, sharper challenge, more time to reflect to identify serious opportunities and issues Threat: confidentiality leaks, over-reliance on AI.	From ad-hoc personal use to a board AI charter with data classification, differentiated analytical tools and skills plan

Three questions to consider for the next 90 days

- 1 Request an AI value dashboard and an agent register
- 2 Adopt a board AI charter and a director skills plan
- 3 Commission an AI-cyber stress test and data-class principles for AI infrastructure

¹ Shadow AI stands for use of non-authorized external AI services.

2. Swiss legal and business frame: no AI Act yet, but duties apply today

Directors of Swiss SMEs cannot wait for an AI statute. Four layers of existing or emerging law already shape what boards must do, and Swiss SMEs are adopting AI quickly but with thin governance.

Layer	What applies (or is coming)	Board implication
Board duties	Art. 716a CO: ultimate direction, organisation and supervision cannot be delegated. Art. 717 CO: duty of care. Art. 754 CO: personal liability for negligent breach of duty causing damage. Commentators read these duties as covering AI and cyber risk and recommend minuting AI decisions ^[28,29] .	Minute AI decisions; show due care.
Data protection	The revised Federal Act on Data Protection (FADP) is technology-neutral and applies to all AI-supported processing of personal data; the FDPIC expects transparency on purpose, functioning and data sources ^[25] .	Know which personal data enters which AI tool.
Cyber reporting	Operators of critical infrastructure must report cyberattacks to the NCSC within 24 hours of discovery (since 1 April 2025; fines since 1 October 2025) ^[26,27] . Banks and investment firms must report within 24 hours and file a full report within 72 hours ^[25] .	Name who reports, and who is on call.
AI-specific rules	No Swiss AI Act yet. The Federal Council signed the Council of Europe AI Convention on 27 March 2025; a consultation draft is due by end-2026, then consultation and Parliament ^[23,24] . The EU AI Act timetable (watermarking 2 Dec 2026; high-risk 2 Dec 2027 and 2 Aug 2028) matters for SMEs serving the EU ^[7] ; confirm reach with counsel.	Watch two clocks: EU and Swiss.

3. Five developments and what they mean for the board

1 The AI value gap: spending doubles, profit impact lags

80% report productivity gains · 37% report any EBIT impact · 6% are “high performers” ^[1,2]

(a) What is the development? McKinsey’s 2026 global survey (1,719 respondents, 97 countries, fielded May–June 2026) finds that 80% say AI improved their individual productivity, but only 37% report a positive EBIT impact and just 6% qualify as “AI high performers” (5%+ of EBIT attributed to AI) ^[1,2]. McKinsey’s conclusion: enterprise value has not kept pace with adoption, because it requires changing how the organisation works, not only deploying tools ^[2]. Meanwhile BCG’s AI Radar 2026 (2,360 executives) shows companies planning to double AI spending, 72% of CEOs now the main AI decision-makers, and 94% prepared to keep or raise spending even if near-term payoffs do not materialise ^[3].

(b) Why it matters: opportunity and threat for the board Opportunity: the 6% show that value is capturable; the difference lies in organisational redesign, not in the technology. Threat: capital is committed on conviction rather than evidence, with the CEO as both sponsor and judge, which exposes the board’s core duty of disciplined capital allocation. The information gap is real: 82% of directors rate the information linking AI outcomes, risks and business performance as fair, poor or not provided ^[14].

Swiss lens 34% of Swiss SMEs consciously use AI (22% in 2024), yet only 34% have set rules on what data employees may enter into AI tools ^[21]. A 2026 study by ti&m and HSLU finds many SMEs acting as followers, with willingness to invest rising despite limited returns so far ^[22]. SME boards can still insist on proof before scale.

(c) What should board members do differently?

- Treat AI as a capital-allocation topic: ask for stage-gated business cases for the largest AI investments (for an SME, the top two or three), each with an EBIT-linked KPI, a baseline and a stop criterion.
- Request a quarterly one-page AI value dashboard: realised financial effect versus plan, full run cost, adoption and incidents.
- Ask in every AI discussion: “Which Workflow or decision is being redesigned, how does our Organization change and Who is responsible for the result?”
- Minute AI investment decisions and their rationale; this documents due care (Art. 717 CO).

2 Agentic AI moves from pilot to operations: software that decides

30%+ of 2026 AI budgets go to agents · EU high-risk rules deferred to Dec 2027, not cancelled [3,4,7]

(a) What is the development? Agentic AI – systems that plan and execute multi-step tasks with limited human involvement – is the main destination of 2026 AI budgets: CEOs have committed over 30% of AI investment to it, and about 90% expect measurable returns from agents this year [3]. Gartner warns that agents create new attack surfaces and that organisations must identify sanctioned and unsanctioned agents [5], reports “ungoverned sprawl” of agents [6], and predicts that by 2030 half of agent deployment failures will stem from insufficient runtime governance [4]. In the EU, the Digital Omnibus (Regulation 2026/1744, in force 27 July 2026) moved high-risk obligations to 2 December 2027 (stand-alone systems) and 2 August 2028 (AI embedded in products); transparency duties under Article 50 applied from 2 August 2026 [7].

(b) Why it matters: opportunity and threat for the board Opportunity: end-to-end process redesign at a scale and speed not possible before; organisations investing in governance tooling were nearly twice as likely to report higher AI value [6]. Threat: decisions taken by software nobody owns, no inventory of what agents can access or do, and liability and reputational exposure when an agent errs. Boards of companies with EU exposure (including Swiss companies whose AI outputs are used in the EU) should plan for the 2026–2028 deadlines now and confirm scope with counsel.

Swiss lens Switzerland has no AI-specific statute yet: the Federal Council signed the Council of Europe AI Convention on 27 March 2025, and a consultation draft is due by end-2026 [23,24]. Existing law applies meanwhile: the FADP governs personal-data processing by agents [25], and Art. 716a and 754 CO govern board oversight [28]. Swiss SMEs selling into the EU should check with counsel whether the EU AI Act reaches them [7].

(c) What should board members do differently?

- Request a simple register of all AI tools and agents in use (sanctioned and shadow): data access, decision rights, owner.
- Approve “autonomy limits”: decisions (e.g. credit, hiring, pricing, external disclosure) that no agent may take without human sign-off.
- Name one accountable executive; require a kill-switch and an AI incident playbook, tested annually.
- If you sell into the EU, ask for an EU AI Act check against the 2026–2028 dates.

3 Mythos-class models: AI changes the economics of cyber risk

7 Apr 2026: frontier model withheld from release · thousands of high-severity flaws reported · 69% of directors rank cyber top AI concern [11,14]

(a) What is the development? On 7 April 2026 Anthropic announced Claude Mythos Preview, a frontier model it chose not to release publicly, deploying it instead through “Project Glasswing” with partners including AWS, Apple, Cisco, CrowdStrike, Google, Microsoft and Palo Alto Networks, backed by up to USD 100m in usage credits [11]. According to the vendor, the model found thousands of high-severity vulnerabilities, including in major operating systems and browsers, and can chain flaws into working exploits with little human guidance [11,13]. These are vendor claims; independent commentary nonetheless concludes that specialised hacking could become a routine, automated process [13].

(b) Why it matters: opportunity and threat for the board Opportunity: the same capability lets defenders find and fix flaws in their own code before attackers do. Threat: the window between discovery and exploitation shrinks, shared software dependencies become exploitable at scale [11], and legacy and operational-technology estates are especially exposed [12]. Regulators are raising expectations: the UK FCA’s March 2026 review expects boards to actively challenge exposure to severe cyber attacks and third-party concentration [12]. 69% of directors name cybersecurity and data privacy as their top AI-related concern [14].

Swiss lens Operators of critical infrastructure (e.g. energy, water, transport) must report cyberattacks to the NCSC within 24 hours of discovery, with fines for non-compliance since 1 October 2025 [26,27]; personal-data breaches are reported to the FDPIC [25]. Directors who ignore known cyber risks can be held personally liable under Art. 754 CO [28].

(c) What should board members do differently?

- Ask your IT lead or managed-security provider for three numbers: time-to-patch for critical vulnerabilities, share of systems scanned with AI-assisted tools, and share of legacy / operational technology assets that cannot be patched.
- Run a board-level tabletop exercise on an AI-enabled attack and on a critical-supplier failure.
- Insist on the fundamentals that advisers rank first: access control, segmentation, zero trust, automated patching [12].
- Confirm whether you must report to the NCSC or FDPIC, who decides, and who is on call within 24 hours.

4 Frontier AI moves on-premises: a trillion-parameter model running on your desk

1 trillion parameters · 748 GB memory · 20 petaFLOPS (FP4) · shipping Q4 2026 [8,9]

(a) What is the development? On 1 June 2026 NVIDIA announced DGX Station for Windows, built on the GB300 Grace Blackwell Ultra Desktop Superchip: 748 GB of coherent memory and 20 petaFLOPS (FP4), able to run AI models of up to one trillion parameters locally and to execute hundreds of agents at once, with Windows security and fleet management [8,9]. Dell, HP, ASUS, GIGABYTE, MSI and Supermicro plan systems for Q4 2026 [9]. At data-centre level, NVIDIA's Vera Rubin platform is in full production [10].

Swiss lens Current online purchasing channels, accessible in Switzerland show Hardware offers ranging from approximately CHF 10'000.- to CHF 100'000.-

(b) Why it matters: opportunity and threat for the board Opportunity: strategic, confidential or regulated data can be analysed in-house without cloud dependency, improving data sovereignty and cost predictability, with workloads scaling to the cloud when needed [8]. Threat: capability once confined to hyperscalers spreads to individual teams outside central IT oversight – shadow AI with frontier power – creating new duties for patching, model provenance and data protection, and new capex-versus-opex trade-offs. The board's question shifts from “which vendor?” to “where does which data run, and who is in control?”

Swiss lens Most SMEs will not buy a trillion-parameter workstation; the practical Swiss choices are public cloud, Swiss-hosted services or open models. Apertus 1.5 (ETH Zurich, EPFL, CSCS; Apache 2.0 licence; released 24 July 2026) is a fully open Swiss model [30], and Swisscom offers Apertus on its Swiss-hosted AI platform [31]. Matching each data class to a hosting option can support FADP compliance.

(c) What should board members do differently?

- Ask management which hosting options (public cloud, Swiss-hosted, open model, own hardware) fit each data class; the board approves the principles.
- Define the data that must never leave company premises or Switzerland, or be processed by external models.
- Require a policy for desk-side (local and authorized) and shadow AI (hidden use of non-authorized external services): approval, security baseline, named owner, asset register.
- Treat AI infrastructure spend (compute, energy, vendor concentration) as a strategic decision, not an IT purchase.

5 Directors adopt AI faster than boards govern it

82–92% of directors use AI for board work · 22% of boards have AI governance processes [15,16,19]

(a) What is the development? 2026 director surveys converge: 82% of US public-company directors used generative AI for board work in the last six months, up from 66% in September 2025 [15]; 92% in OnBoard's sample of 531 governance professionals [16]. Governance lags: only 22% report processes guiding that use [19], most boards lack board-specific AI policies [17], and 49% of directors have heard of colleagues using consumer AI tools rather than company-approved platforms [15]. Skills lag too: 71% say the board must strengthen its AI skills [14], and only about a third believe the board has the right skills to oversee AI, cybersecurity and data governance [18].

(b) Why it matters: opportunity and threat for the board Opportunity: AI can digest voluminous board papers, surface insights from data and sharpen challenge of management; directors are beginning to use it as a strategic partner [20]. Threat: confidential board material leaking into consumer tools, over-reliance weakening human judgment [14], and unclear accountability. Recruitment habits also hold boards back: 81% rate cultural fit as very important, but only 27% rate specialised expertise such as AI or cybersecurity the same [14].

Swiss lens SME boards are small and often owner-led, so a one-page AI charter, minuted AI decisions (Art. 716a and 717 CO) and one named AI sponsor on the board are proportionate [28,29]. Only 34% of Swiss SMEs have clear rules on what data may enter AI tools [21].

(c) What should board members do differently?

- Adopt a one-page board AI charter within 90 days: approved tools only, no confidential papers in consumer AI, human verification of AI-generated minutes and summaries.
- Inventory what directors already use, then offer a secure, approved alternative.
- Add AI and cyber fluency to the skills matrix, induction and board evaluation; set a minimum annual learning programme.
- Use AI as a devil's advocate on management proposals; accountability stays with directors.

4. Starting points for the sharing of experiences

1. How do you use AI when preparing your Board meetings? What are your experiences?
2. What rules govern how we use AI as a board?
3. Which data would we never let leave our premises, and who decides?
4. Do you know which agents exist in your company and what they may decide alone?
5. How fast do we patch critical vulnerabilities, and when did the board last rehearse an AI-enabled attack?
6. Which of our AI investments can we tie to EBIT today, and which are acts of faith?

5. Sources

- [1] McKinsey & Company, "The state of AI in 2026: On the road to ROI", 25 Aug 2026 (survey of 1,719 participants, 97 countries, 4 May–8 Jun 2026). mckinsey.com
- [2] McKinsey & Company, "AI is changing work. Now it has to change the organization", 2026. mckinsey.com
- [3] Boston Consulting Group, "AI Radar 2026: As AI Investments Surge, CEOs Take the Lead", 15 Jan 2026 (2,360 executives, 640 CEOs). bcg.com. The 94% figure is as reported in press coverage of the BCG study.
- [4] Gartner, "Gartner Announces Top Predictions for Data and Analytics in 2026", press release, 11 Mar 2026. gartner.com
- [5] Gartner cybersecurity predictions 2026 (A. Michaels), as reported in "Gartner: Tighten Up AI Governance or Face the Consequences", [Tech.co](https://tech.co), 2026.
- [6] Gartner Digital Workplace Summit (M. Goss), as reported in "Govern your bots carefully or chaos could ensue", [The Register](https://the-register.com), 30 Apr 2026.
- [7] Cloud Security Alliance Labs, "EU AI Act's High-Risk Deadline: Deferred, Not Cancelled", 2026 (on Regulation (EU) 2026/1744, Digital Omnibus on AI, in force 27 Jul 2026).
- [8] NVIDIA Newsroom, "NVIDIA DGX Station for Windows Puts a Trillion-Parameter AI Supercomputer on Every Enterprise Desk", 1 Jun 2026 (GTC Taipei). nvidianews.nvidia.com
- [9] Enterprise DNA, "NVIDIA Puts a Trillion-Parameter Supercomputer on Your Desk", 2 Jun 2026 (hardware specifications and partner availability, Q4 2026).
- [10] Techzine Global, "Nvidia's plans for Vera Rubin, DGX Station, agents and robots", Jun 2026.
- [11] Bloomsbury Intelligence and Security Institute (BISI), "Claude Mythos and the Acceleration of Cybersecurity Risk", 2026. bisi.org.uk
- [12] Fladgate LLP, "Project Glasswing: what company boards need to know about Claude Mythos", 2026 (incl. UK FCA March 2026 review).
- [13] The Conversation, "Claude Mythos and Project Glasswing: why an AI superhacker has the tech world on alert", 2026.
- [14] PwC, "2026 Annual Corporate Directors Survey", Sep 2026 (n = 561). pwc.com
- [15] Diligent, "Board AI use rises: Director Confidence Index", Jun 2026 (104 US public-company directors). diligent.com
- [16] OnBoard, "2026 Board Effectiveness Survey", sixth annual report, closed 1 Jun 2026 (531 governance professionals).
- [17] Deloitte Center for Board Effectiveness, "How boards are using AI today", Board Practices Quarterly, Jul 2026 (Society for Corporate Governance members).
- [18] BDO, "2026 BDO Board Survey: Governance Under Pressure", 2026 (200+ US public-company directors).
- [19] Diligent Institute / Corporate Board Member, "What Directors Think 2026".
- [20] Harvard Law School Forum on Corporate Governance, "2026 Corporate Governance Trends to Watch" (PwC), 8 Feb 2026. corpgov.law.harvard.edu
- [21] Swiss SME Portal (SECO), "AI gains ground among Swiss SMEs", 5 Nov 2025 (AXA/Sotomo labour-market study). kmu.admin.ch
- [22] ti&m and Lucerne University of Applied Sciences and Arts (HSLU), "AI Maturity Study 2026", presented 5 Feb 2026.
- [23] Swiss Federal Chancellery, "Regulation of AI" (Federal Council decision of 12 Feb 2025; consultation draft due by end of 2026). bk.admin.ch
- [24] Chambers Practice Guides, "Artificial Intelligence 2026: Switzerland" (Council of Europe AI Convention signed 27 Mar 2025).
- [25] Chambers Practice Guides, "Cybersecurity 2026: Switzerland" (FADP and AI, FDPIC, NCSC reporting, banks).
- [26] Swiss Federal Council, "Reporting cyberattacks on critical infrastructure mandatory from 1 April", 7 Mar 2025. admin.ch
- [27] Baker McKenzie, Connect On Tech, "Switzerland: Reporting cyberattacks on critical infrastructure mandatory from 1 April" (fines from 1 Oct 2025).
- [28] Bratschi, "Cybersecurity und Haftung in der Schweiz" (Art. 716a, 717, 754 CO). bratschi.ch
- [29] Swiss Board Forum, "Künstliche Intelligenz und persönliche Haftung: Strategische Pflichten für den Verwaltungsrat". swissboardforum.ch
- [30] ETH Zurich, EPFL and CSCS, "Apertus 1.5: Building the next generation of open AI infrastructure", 24 Jul 2026. cscs.ch

[31] EPFL / GGBA, “Switzerland launches Apertus, a fully open multilingual large language model” (Swisscom sovereign AI platform).

[32] [Digitec](#) and [Galaxus](#) online offers, dated October 13, 2026

Method and caveats. Sources are 2026 publications by consulting firms (McKinsey, BCG, PwC, Deloitte, BDO), Gartner, NVIDIA, the Harvard Law School Forum on Corporate Governance, Swiss federal bodies, law firms and universities. Global and director surveys mostly cover large and US companies and some are sponsored by board-software vendors; treat their figures as directional for Swiss SMEs. Swiss SME statistics come from the AXA/Sotomo study (via SECO) and the ti&m/HSLU study. Mythos-related capability claims originate from the vendor (Anthropic, the maker of Claude) and are corroborated only in part by independent commentary. Figures marked “as reported” come from press coverage of the primary study. Swiss and EU legal points summarise public sources, are not legal advice, and should be confirmed with counsel.