

**2026**  
**LUXEMBOURG**  
**GRC SUMMIT**  
powered by NC3

**HOW SHOULD CYBER GOVERNANCE  
EVOLVE IN THE AGE OF AI?**

**OUTCOME PAPER**

Key insights and recommendations from the forum on the evolution of cybersecurity governance in the age of AI, held during the Luxembourg GRC Summit 2026, exploring the implications for trust, accountability and risk management.

# Table of contents

---

|    |                                                     |
|----|-----------------------------------------------------|
| 3  | Summary                                             |
| 4  | Rethinking Trust and Accountability                 |
| 5  | How AI is Reshaping the Cyber Risk Landscape        |
| 6  | Bridging Organisational and Product Risk Management |
| 7  | Preparing Practitioners                             |
| 8  | Meaningful Human Oversight and Accountability       |
| 9  | Key Takeaways                                       |
| 10 | Recommendations and Action Points                   |



# HOW SHOULD CYBER GOVERNANCE EVOLVE IN THE AGE OF AI?

## Summary

---

The current report provides an account of the discussions and outcomes of the forum “How should Cyber Governance evolve in the age of AI?”, held during the second edition of the GRC Summit in Luxembourg on 3<sup>rd</sup> and 4<sup>th</sup> of June 2026. The Summit brought together Cybersecurity Governance, Risk and Compliance (GRC) professionals, practitioners and decision-makers to exchange knowledge, examine emerging threats, and explore evolving practices in the field.

Discussions explored how trust, transparency, and accountability should be understood in an environment where AI systems can introduce uncertainty and less predictable outcomes. Participants considered whether traditional cybersecurity and governance frameworks, largely designed around deterministic systems, remain fit for AI-enabled environments and how risk management, assurance, and audit appr-

roaches may need to adapt. The role of emerging regulation, including the EU AI Act, was also discussed as a potential framework for addressing new risks and clarifying responsibilities. This was complemented by discussions on the internal governance of cybersecurity and GRC functions, including evolving roles and oversight mechanisms.

Finally, participants highlighted the need for GRC and cybersecurity professionals to develop new skills and competencies to effectively govern AI related risks and navigate an increasingly complex technology and regulatory landscape.

# Rethinking Trust and Accountability

# 01

The forum begins with a fundamental question: in an era profoundly reshaped by artificial intelligence, how should we understand, establish, and maintain trust? If cybersecurity is ultimately about creating and maintaining trust, then the age of AI requires us to reconsider what trust means in the first place. In a world where systems can learn, generate, recommend, and increasingly act with varying degrees of autonomy, trust can no longer be treated simply as a binary decision: secure or insecure, trusted or untrusted. Instead, trust may need to be understood as a spectrum. Crucially, trust can also no longer be a one-time judgement made at the point of deployment. As AI systems, models, data, threats and operating environments evolve, trust must be continuously tested, reassessed and maintained throughout the system's lifecycle.


This approach shifts the focus of cybersecurity governance from asking "Can we trust this system?" to asking:

***"Under what conditions, to what degree, and for which purposes can we trust it, and how do we ensure that this trust remains justified over time?"***

Transparency and responsibility become essential to answering these questions:

## CYBER TRUST WITHIN ORGANISATIONS IN THE AI ERA:

### WHAT MATTERS MOST?



*Responses from forum participants to the question "What matters most for cyber trust within the organisation in the AI era?" (63 respondents)*

organisations must be able to make the boundaries, risks and behaviour of AI systems sufficiently visible, while ensuring that responsibility for their use remains clearly assigned. In the age of AI, therefore, trust is not a destination or a binary status; it is a continuous governance process.

Accountability presents a further challenge as AI becomes increasingly embedded in organisational processes and capable of making or influencing decisions. Responsibility may be distributed across multiple actors, including AI providers, system integrators, the organisation deploying the system, business owners, security

and risk functions, and senior management. This challenge becomes even more pronounced in multi-agent AI environments, where several autonomous or semi-autonomous agents may interact, influence one another, and contribute to a final decision or action. Governance must therefore address accountability both across the organisational actors involved and within increasingly complex AI architectures themselves.

Taken together, these challenges suggest that the central issue is not whether

existing governance principles remain relevant, but whether organisations can make those principles operational in an increasingly dynamic and interconnected AI environment. Governance must provide mechanisms to assess and reassess trust over time, establish clear accountability across complex systems and actors, and adapt controls as AI capabilities, dependencies and risks evolve.

## How AI is Reshaping The Cyber Risk Landscape

# 02

Traditional cybersecurity risk analysis has largely been built around relatively deterministic systems, where defined algorithms, logic, and configurations produce behaviour that can be tested, modelled, and anticipated with a high degree of confidence. AI challenges this foundation. AI systems are inherently stochastic; their outputs can vary depending on the data they process, the context in which they operate, the interactions they receive and, in some cases, the probabilistic behaviour of the underlying models. Moreover, AI systems and the ecosystems around them can evolve throughout their lifecycle, as models, data, applications, dependencies

and operating conditions change. This introduces a fundamentally more dynamic dimension to cybersecurity risk. The risk profile of an AI-enabled system can therefore change not only when its configuration changes, but also as its inputs, context, capabilities and environment evolve.

*The underlying principles of risk management therefore remain relevant, but their application must evolve to account for a risk environment characterised by greater speed, uncertainty, autonomy, and dependence on external systems, models and providers.*

This is also where some of the main governance challenges emerge. Overcoming these challenges requires a shift from point-in-time risk assessment to continuous risk monitoring and reassessment. Rather than treating risk as a relatively stable attribute established before deployment, organisations need to recognise that the risk profile of their operational environment can change throughout its lifecycle. Changes in models, data, system integrations, user behaviour, threat activity or operating context can alter both the likelihood and potential impact of adverse outcomes.

Discussion around privacy-related risk emphasised the fact that privacy deserves to be treated as more than an extension of traditional data-protection controls. The ubiquity of AI systems has considerably increased privacy risks, as these systems can derive information, make inferences and combine datasets in ways that were not anticipated. The question is no longer only whether data is appropriately protected, but also what

an AI system can infer, reconstruct or expose from the data to which it has access. This becomes particularly important as AI systems are increasingly integrated into enterprise environments and can access and combine information from multiple sources.

The governance challenge is therefore to bring data governance, privacy and information security much closer together, ensuring that organisations consider not only how data is protected, but also how AI can transform what can be derived from it. Risk management also needs to become more dynamic, with mechanisms to continuously monitor changes, test system behaviour, reassess risk exposure and adjust controls as circumstances evolve. The existing governance frameworks also need to evolve accordingly, extending their scope beyond the security of traditional IT assets to address the security, resilience and control of AI-enabled capabilities and the dependencies they rely on.

## Bridging Organisational and Product Risk Management

# 03

Traditional cybersecurity risk management has generally looked at risk from the perspective of the organisation. Risks are assessed at enterprise, business unit, portfolio or programme

level and evaluated against an organisation's risk appetite and tolerance thresholds. Typical considerations include financial or reputational loss, compliance exposure and operational

disruption. This approach remains relevant. Some EU frameworks explicitly follow it: DORA, for example, requires financial entities to define their ICT risk tolerance in accordance with their risk appetite.

AI and recent EU product legislations introduce another perspective linked to the product itself. The AI Act reflects this approach for AI systems: risks must be assessed in light of the system's intended purpose and reasonably foreseeable misuse.

The Cyber Resilience Act follows a similar logic for products with digital elements. Product-oriented assessments ask different questions:

*Can the system fulfil its intended function safely and reliably, what are its*

*foreseeable failure modes, how could it be misused, who could be harmed, and how could a cybersecurity issue affect the product's function?*

Organisational risk management and product risk management are compatible, and should even be considered as complementary.

Finally, intended purpose is also particularly relevant before an AI system is introduced. The starting point should be the business objective rather than the availability or popularity of AI. Organisations should consider what problem needs to be solved, what role AI would play, and whether its expected benefits justify the costs and risks compared with other solutions.

## Preparing Practitioners

# 04

Cybersecurity professionals will need to develop new competencies to govern AI effectively. This includes knowledge of AI risk management and regulatory requirements, but also more operational skills such as AI system documentation, assurance, data governance, supply chain and third-party risk, human oversight, logging and monitoring, and the specific risks associated with agentic AI. Business process analysis and the

ability to communicate AI-related risks clearly to management and boards will also become increasingly relevant.

For SMEs in particular, AI governance should be translated into practical steps rather than remain at the level of general principles. Useful support could include simple checklists, questions to ask AI providers, guidance on defining the intended purpose of a system, clarity on

how data is used, basic requirements for human oversight, practical risk assessment templates and clear indications of what evidence should be retained.

AI literacy also needs to extend beyond specialists. If employees and other users are expected to supervise AI systems,


challenge their outputs or decide when human intervention is required, they need a basic understanding of what these systems can and cannot do, where their limitations lie and how they can be used safely.

## Meaningful Human Oversight and Accountability

# 05

A human ultimately needs to remain responsible and accountable for the use of AI. However, simply inserting a person into a process does not automatically make it safe. Different levels of human involvement may be needed depending on the level of risk and the degree of trust that can reasonably be placed in the system. This can range from a person actively approving or contributing to decisions, through systems operating under human supervision, to situations where humans remain largely outside the operational loop.

For human oversight to be meaningful, the person responsible needs the ability to understand what the system is doing, challenge its outputs, intervene when necessary, escalate concerns, override decisions or stop the system. Existing controls such as segregation of duties, process gates and checkpoints can also help maintain effective human control.



Responses from Forum participants to the question "What is currently the least understood enterprise AI cyber risk?" (45 respondents)

Automation bias makes oversight more difficult in practice: when an AI system repeatedly produces acceptable results, users may stop checking its outputs with the same level of attention. This concern is also reinforced by the risk of collective unlearning. As people rely more heavily

on AI systems, they may gradually lose the knowledge and contextual understanding needed to judge whether an output is appropriate. Paradoxically, the more reliable the system appears, the easier it can become for human oversight to turn into a formality, thus weakening the very human oversight on

which many governance and regulatory approaches depend.

***Meaningful human oversight should be designed as an operational control, with clear roles, intervention points and escalation mechanisms, rather than treated as a procedural requirement.***

## Key Takeaways

---

The discussion showed that AI governance is moving away from static, document-based approaches towards continuous governance lifecycle risk management and evidence-based trust. Trust can no longer be treated as a one-time decision made before deployment. It needs to be demonstrated through monitoring, testing, documentation, transparency, and the ability to reassess risk as systems, data, users, and dependencies evolve. This also means looking beyond the model itself and governing the wider ecosystem in which AI operates.

At the same time, human oversight, accountability and control remain central. As AI systems become more autonomous and agentic, organisations need to define who is responsible, what authority AI systems should have, when human intervention is required, and how that intervention can be exercised in

practice. Oversight must remain meaningful, supported by sufficient competence, attention and authority. Without this, automation bias, overdependence and gradual loss of expertise risk weakening the very controls that governance frameworks rely on.

More broadly, effective risk management in today's landscape will require cross-functional coordination, practical implementation and new competencies. Legal obligations, technical controls, standards and operational risk management need to be connected rather than handled separately. Cybersecurity and GRC professionals will need to expand their skills, while SMEs and other organisations need guidance that is concrete enough to apply in practice.

# Recommendations and Action Points

---

**01.** Move from one-off assessments to **continuous governance**. Establish monitoring, review, and reassessment mechanisms that continue after deployment.

**02.** Build an **evidence base** for AI trust assurance. Maintain documentation that can demonstrate **how the system is governed** and **how it performs** over time.

**03.** Design **human oversight** as an operational control. Do not treat it as a formality.

**04.** Apply **least privilege** and **least agency** to AI agents. **Limit agents** to minimum permissions, tools, and actions required for their intended function.

**05. Govern the full ecosystem**, not only the model. Extend risk assessment to **third-party dependencies**.

**06.** Establish **clear cross-functional accountability**. Define **responsibilities** across cybersecurity, legal, data management, procurement, business functions, and senior management.

**07. Map regulatory requirements** to technical and operational controls.

**08.** Provide **practical AI governance tools** for SMEs and less mature organisations.

**09.** Invest in **AI governance skills** and preserve **human expertise**.

**10.** Seek expert and consultant **guidance** where expertise is needed.

Thank you for reading. We hope these insights inform your work and encourage further dialogue on cybersecurity governance.

Thank you to our participants, Advisory Board and partners for sharing their expertise, **and to you for taking these discussions forward.**



**nc3.lu**

National Cybersecurity  
Competence Center  
**LUXEMBOURG**

**IN COLLABORATION WITH**



**LHC**

Luxembourg House  
of Cybersecurity



**LUXEMBOURG  
CYBERSECURITY  
FACTORY**