Nemko Digital

Scaling AI with a Risk-Based Control Framework Webinar Report

---

Executive Summary

Nemko Digital's webinar, "Scaling AI with a Risk-Based Control Framework: From Experimentation to Enterprise Governance," examined how organizations can move from promising AI pilots to reliable, production-ready AI systems by adopting a structured, risk-based approach to control design and governance. The central message was clear: regulatory compliance matters, but it is not sufficient on its own. Organizations seeking to scale AI responsibly must evaluate both the societal risks that regulators focus on and the organization-specific risks that determine whether AI can be trusted, governed, and operationalized in day-to-day business processes.

The discussion combined regulatory context, standardization developments, business-oriented impact assessment, and a practical government-sector case study. Alicja Halbryt introduced the global AI risk landscape and explained why AI risk cannot be reduced solely to legal compliance. Pep van der Laan then translated those concepts into an applied example, showing how a government body used a structured AI risk and control framework to move a generative AI use case from a stalled pilot toward safe production deployment.

Two Perspectives on AI Risk

Alicja Halbryt then introduced the wider AI regulatory landscape. She explained that different jurisdictions are taking different approaches to AI governance. The United States and the United Kingdom were presented as examples of more pro-innovation approaches, while South Korea emphasizes human-centric protections, China places significant emphasis on state oversight and enforcement,

and Singapore's framework highlights risk assessment, human accountability, technical controls, and end-user responsibility.

The EU AI Act was discussed as a leading example of a risk-based regulatory approach. Alicja emphasized that regulations primarily focus on societal risks, such as harms to health, safety, and fundamental rights. The deck illustrated the AI Act's categories through examples: prohibited AI includes social scoring, predictive policing, and emotion recognition in workplaces and education; high-risk AI includes biometric identification, employee performance tracking, and credit scoring; limited-risk AI includes chatbots, virtual assistants, and AI-generated media; low-risk AI includes spam filters, grammar checkers, and video-game AI; and general-purpose AI models are subject to GPAI-specific requirements.

AI Risk Category Discussed	Examples Mentioned	Type of Obligation
Prohibited AI	Social scoring, predictive policing, emotion recognition in workplaces and education	Prohibited
High-risk AI	Biometric identification, employee performance tracking, credit scoring	Conformity assessment and strict controls
Limited-risk AI	AI-written articles, chatbots, virtual assistants, generative AI media	Transparency requirements

Nemko Digital

Low-risk AI	Spam filters, grammar and spelling tools, video-game AI	No mandatory obligations; subject to codes of conduct
General-purpose AI	GPT-4, DALL-E, BERT, LLaMA	GPAI-specific requirements

The session also covered European standardization efforts intended to support AI Act implementation. Alicja noted that the European Commission requested standards from European standardization bodies, including CEN and CENELEC, and that the process is complex because it requires alignment among many experts across multiple technical domains. The deck indicated that full coverage of standardization requirements is not expected before 2027. Among the more advanced standards discussed were prEN 18228 “AI Risk Management System,” prEN 18286 “Quality Management System,” and prEN 18282 “Cybersecurity of AI system.”

A key insight from this segment was that regulatory and standardization frameworks are necessary, but they do not fully answer the question of how an individual organization should manage AI adoption. Alicja used ISO/IEC 42005:2024 on AI system impact assessment to explain that a complete impact assessment should consider not only impacts on individuals and society, but also impacts on the organization itself.

From Regulatory Risk to Business Risk

The webinar’s conceptual turning point was the distinction between system-level risk and organization-level risk. Alicja explained that when people think about AI risk, they often begin with the impact on individuals and society. This is appropriate, because these are the risks regulators are usually designed to address. However, organizations also need to examine how potential AI-related events affect their own

operations, decision quality, accountability, reputation, customer service, and strategic objectives.

This dual view requires assessing both impact and likelihood. A risk is not equally severe in every context. The same conceptual AI risk may be urgent in one use case and irrelevant in another. The speakers therefore emphasized that risk management must be connected to the specific AI application, the organization's existing controls, the intended users, the lifecycle stage, and the organization's risk appetite.

Perspective	Primary Concern	Typical Questions
System or societal perspective	Impact on individuals, society, health, safety, and fundamental rights	Could the AI system harm people, undermine fairness, reduce transparency, or affect legally protected interests?
Organization-level perspective	Impact on business processes, accountability, trust, efficiency, and operational risk	Could the AI system undermine service quality, create unowned decisions, damage stakeholder trust, or introduce unmanaged dependencies?
Integrated risk perspective	Severity based on both impact and likelihood	Which risks are relevant, how likely are they, how severe would they be, and what controls are needed?

To make this interactive, the speakers ran a short quiz asking attendees where they were spending their time when it came to AI risk: organization level, system level,

or a balance of both. The disclosed results showed that 21% were almost entirely focused at the organization level, 26% reported a balanced focus, 5% were mostly at the system level, and 11% were almost entirely at the system level. The largest group was described verbally as being mostly focused at the organization level. Pep connected these results to the case study, noting that they reflected the same pattern observed in practice: much of the value in AI governance is created by addressing organizational readiness, not only regulatory readiness.

Case Study — Generative AI in Government

The most detailed part of the webinar was a government-sector case study involving a generative AI use case. Pep explained that large corporations can ask a government body for a pre-assessment that provides upfront clarity on how laws and regulations apply in specific scenarios. Such pre-assessments support internal decision-making, compliance assurance, and the avoidance of unexpected regulatory outcomes.

For the government body, the process was slow and time-consuming because it depended on many rules, regulations, formal requirements, prior statements, and jurisprudence. The organization developed an AI solution to support human assessors in two main ways. First, the system could assess whether formal requirements were met and draft a report on whether a request could legally be considered. Second, it could help find relevant clauses in regulations, jurisprudence, and supporting documentation so that human assessors could work faster and with more structured information.

The case began with a promising pilot and enthusiastic frontline workers. However, the pilot did not automatically translate into production deployment. Pep described several hurdles: the organization lacked sufficient AI expertise to identify and delimit AI risks, decision makers were risk-averse, there was no shared perspective on how to assess AI risks, and no clear ownership existed for AI-related approval decisions. As a result, go-live decisions stalled, and the organization could not yet realize the expected benefits in turnaround speed and operational efficiency.

Nemko Digital

Scaling Hurdle	Practical Effect
Promising pilot but no clear path to production	The AI use case demonstrated value but could not move into stable operational use.
Insufficient AI expertise across the organization	Stakeholders struggled to identify which AI risks were real, relevant, or severe.
Risk-averse decision-making	Novel AI applications were difficult to approve, even when the use case was bounded.
No shared risk assessment method	Stakeholders lacked a common language for judging risk and controls.
Missing ownership	Product teams were sent from stakeholder to stakeholder without a clear decision route.
Stalled go-live decisions	The organization and its clients could not benefit from faster and more efficient processes.

The Risk-Based Control Framework Applied in the Case Study

Nemko Digital's approach began with a structured process: kickoff, document analysis, risk atlas development, risk analysis, control identification, framework completion, and final reporting and handover. The key was to move the conversation from general AI anxiety to a concrete analysis of the specific use case.

Nemko Digital

The team built a client-specific version of a risk atlas, drawing on IBM's AI Risk Atlas and adapting it to the client context. Pep emphasized that a risk atlas helps organizations create a more complete view of potential AI-related risks. It also prevents a fragmented process where every stakeholder adds risks from their own perspective without a shared structure. In the case study, the full risk atlas contained 87 generally described AI risks. After evaluating relevance to the specific use cases, the team determined that 40 of the 87 risks were applicable.

A crucial methodological distinction was the move from conceptual or gross risk to net risk. Pep explained that organizations often overestimate risk by considering everything that could theoretically go wrong with AI in general. A useful risk evaluation must instead examine the actual use case, its scope, and the controls already in place. Many organizations already have policies, access controls, data security measures, privacy safeguards, and operational procedures. These existing controls affect the real remaining risk. The remaining risk, or net risk, is what must be managed through additional controls and ultimately accepted by the right accountable owner.

Risk Concept	Meaning in the Webinar	Why It Matters
Conceptual risk	A risk that may exist for AI systems in general	Useful for completeness, but too broad for operational decisions.
Gross risk	Theoretical risk before considering controls	Can exaggerate risk if organizational context is ignored.
Net risk	Use-case-specific risk after considering scope	The practical basis for prioritizing mitigation.

	and existing controls	
Residual risk	Remaining risk after additional controls are implemented	Must be formally accepted by an accountable owner.

The risk assessment used likelihood and impact to prioritize action. High risks were the primary focus, because the objective was not to eliminate every theoretical risk but to bring significant risks down to a level consistent with the organization's risk appetite. Pep described the goal as bringing high risks down to a medium or acceptable level so that the organization could deploy use cases that were secure, reliable, trustworthy, and "good enough" for production.

Control Taxonomy: Strategy, Prevention, Detection, and Correction

After identifying and prioritizing risks, the next step was to define controls. The webinar presented a four-pillar taxonomy: strategy, prevention, detection, and correction. Strategy focused on governance. Prevention focused on data and design. Detection focused on technical defence, testing, and monitoring. Correction focused on human oversight and validation.

The deck listed numerous control focus areas, including user instructions, testing, monitoring, human validation, architecture guidelines and patterns, documentation, work instructions, LLM guidelines, acceptance of residual risk, the four-eyes principle, access control, data quality processes, data masking, internal reporting options, user notification, logs and data retention policy, bias testing, change management, and DPIA pre-scans.

Control Pillar	Focus Area	Examples Discussed or Shown
Strategy	Governance	Ownership, decision rights, residual risk acceptance, AI Board escalation.
Prevention	Data and design	Architecture patterns, data quality processes, data masking, LLM guidelines, work instructions.
Detection	Technical defence	Testing, evaluation, validation, verification, monitoring, logs, retention policies, bias testing.
Correction	Human oversight	Human validation, four-eyes principle, internal reporting, user notification, corrective feedback loops.

Pep emphasized a defence-in-depth logic: multiple controls may be needed for a single risk, and controls can operate at different points in the lifecycle. The case study identified more than 150 potential controls. However, after prioritization, only about one-third of the risks required additional measures, and the team ultimately identified 39 new or updated controls that needed effort to bring the AI use case within an acceptable risk profile.

Nemko Digital

Governance Flow, Ownership, and Decision-Making

The speakers made clear that a control framework cannot remain a one-off assessment document. To enable scaling, risk management must be integrated into the AI lifecycle. The deck showed the relationship between an AI governance flow and an AI development flow. Governance activities included impact assessment, use-case approval, use-case onboarding, data-access approval, risk evaluation, model approval, and risk and value monitoring. Development activities included use-case creation, data access and preparation, feature engineering, model training and tuning, model registration, model deployment, model serving, and monitoring.

Ownership was a major theme. The case study defined a risk management playbook and ownership principles. Accountability for individual controls depends on the nature of the control. For example, model monitoring may sit with a product technology lead, while AI literacy among frontline workers may be owned by a business team lead. Accountability for risk acceptance generally lies with the business owner, while oversight and coordination sit with the product owner.

Pep also discussed the role of an AI Board where decision norms are missing. Such a board can bring together technology, security, data, legal, ethics, risk, and business stakeholders. This cross-functional structure helps organizations avoid treating AI governance as a purely technical exercise. Instead, it creates a forum where different perspectives can be weighed and where “good enough” can be defined for the organization in a transparent and consistent way.

Governance Element	Purpose
Impact assessment	Establishes the relevance and severity of AI-related impacts.
Use-case approval	Determines whether an AI use case should proceed.

Data-access approval	Ensures appropriate data governance and access controls.
Risk evaluation	Prioritizes relevant net risks based on likelihood and impact.
Model approval	Confirms readiness before deployment or production use.
Risk and value monitoring	Tracks whether the system remains safe, useful, and aligned after deployment.
AI Board	Provides cross-functional decision-making where ownership or norms are unclear.

Q&A and Discussion Highlights

The formal Q&A did not include a substantial set of audience questions in the available transcript. Instead, the closing discussion evolved into a moderated exchange between Pep and Alicja, using the case study to draw out additional lessons. This exchange functioned as a reflective Q&A on the practical experience of implementing the risk framework.

Alicja emphasized that the risk atlas was valuable because it gave the client a reality check. It showed risks that some stakeholders had not previously recognized, but it also reassured the organization that the list of risks was not infinite. In other words, the framework made risk visible and manageable rather than abstract and overwhelming.

She also highlighted the importance of communication between different professional groups. In the case study, ethicists, engineers, and business stakeholders needed to have a productive conversation about topics such as data bias, data privacy, value alignment, and operational feasibility. Pep added that business users were eager to adopt the technology because they could see the practical benefits, while engineers had to make the solution work and compliance or ethics stakeholders had to ensure that adoption remained responsible.

A particularly important insight from the discussion was that privacy was not necessarily the dominant concern, even in a government setting. Alicja noted that the most significant issues in this case related to governance and value alignment, including transparency, testing, and the risk of over-reliance or under-reliance on generative AI outputs. This was a valuable reminder that AI risk profiles are use-case-specific. Organizations should not assume in advance that one category, such as privacy, will always be the highest priority.

Discussion Question or Theme	Key Takeaway
What was most striking when entering the organization?	The risk atlas provided both a reality check and reassurance: many risks existed, but they could be structured and controlled.
What communication challenges emerged?	Ethicists, engineers, and business stakeholders needed a shared language to discuss bias, privacy, feasibility, and value.
Which controls were most important?	Governance and value alignment were more significant than expected; privacy was not the only or primary concern.

Nemko Digital

What was distinctive about the generative AI case?	Over-reliance and under-reliance on AI outputs were important risks that needed control through training, oversight, and process design.

Key Takeaways

- **AI scaling is not only a technical challenge:** A successful pilot does not automatically mean an AI system is ready for production. Organizations need governance structures, approval routes, risk ownership, and ongoing monitoring to scale AI responsibly.
- **Regulation matters, but it is not enough:** The webinar highlights that regulations primarily focus on societal risks, including health, safety, and fundamental rights. Businesses must also evaluate how AI affects their own processes, decisions, reputation, accountability, and service quality.
- **Risk management should be use-case-specific:** Not every theoretical AI risk applies to every AI system. A practical framework helps organizations identify which risks are relevant, assess their likelihood and impact, and prioritize the controls that matter most.
- **Net risk is more useful than theoretical risk:** Organizations often already have privacy, security, access, data, or operational controls in place. Effective AI risk assessment considers these existing controls and focuses on the remaining risk that still needs to be managed.
- **Controls need to cover governance, prevention, detection, and correction:** Responsible AI requires more than technical testing. It also depends on governance measures, data and design practices, monitoring, human oversight, user instructions, documentation, AI literacy, and clear residual-risk acceptance.
- **Clear ownership prevents stalled AI deployment:** The case study showed that AI initiatives can become stuck when no one has the mandate to decide whether a system is “good enough” for production. Defined ownership principles and cross-functional decision-making help organizations move forward with confidence.

Nemko Digital

- **Responsible AI governance can accelerate innovation.** A structured risk-based framework does not slow AI adoption; it creates the trust, clarity, and accountability needed to scale AI use cases safely and effectively.

Overall Value and Recommended Next Steps

The overall value of the webinar lay in its practical framing of AI governance as an enabler of scaling, not merely a compliance burden. By combining regulatory awareness, impact assessment, risk prioritization, control design, and lifecycle governance, the speakers showed how organizations can move beyond abstract AI concerns and toward responsible deployment. The government case study was especially useful because it demonstrated that even risk-averse organizations can create a path to production when they have a clear framework, shared language, and accountable decision-making structure.

For attendees, the most immediate next step is to assess whether their organization has a comparable AI risk and control framework in place. Organizations should begin by inventorying AI use cases, mapping relevant risks, identifying existing controls, and clarifying who owns control implementation, risk acceptance, and ongoing monitoring. They should also evaluate whether their governance process covers the full AI lifecycle, from use-case ideation and data access through deployment, serving, monitoring, and feedback.

Recommended Next Step	Intended Outcome

Create or update an AI use-case inventory	Establish visibility over where AI is being used or planned.
Apply a structured AI risk atlas	Ensure risks are identified comprehensively and consistently.
Evaluate net risk, not only theoretical risk	Focus resources on the risks that remain after scope and existing controls are considered.
Define control ownership and risk acceptance	Prevent stalled decisions and unclear accountability.
Integrate governance into the AI lifecycle	Ensure risk management continues after deployment.
Establish or strengthen a cross-functional AI Board	Enable timely decisions across technology, legal, risk, ethics, security, data, and business functions.

In conclusion, responsible AI scaling depends on turning trust into an operational capability. Nemko Digital's webinar demonstrated that this requires more than awareness of regulations. It requires a practical control framework, clear governance flows, and an organizational commitment to making AI systems trustworthy, accountable, and useful at scale.

Nemko Digital

Special Offer: Call with our AI Act Experts Now!

Don't wait for the next step.

As a special offer for webinar participants, we are providing an exclusive opportunity to discuss your specific AI compliance challenges directly with our experts.

- Open to Webinar Participants
- Complete Application Form
- Share the topic you want to talk about

Ready to get started? Scan the QR code in the webinar materials or visit the link below to apply for your expert consultation:

CLICK HERE: https://digital.nemko.com/scaling-ai-risk-based-control-framework-offer

Nemko Digital

Nemko Digital

Call with our AI Act experts now!

Don’twaitforthenext step

- Open to Webinar Participants
- Complete Application Form
- Share topic you want to talk about

https://links.nemko.com/scaling-ai-offer

Scan to apply