Nemko Digital

Nemko Digital Webinar Report - AI Governance Tools & Technologies: Building Trust and Compliance in 2025

Executive Summary

Nemko Digital hosted an insightful webinar titled "Scaling AI, responsibly: Tools & Technology deep dive." The session featured two of Nemko's leading experts, Bas Overtoom, Senior AI Trust Expert, and Pepijn (Pep) van der Laan, Global Technical Director. They provided a comprehensive overview of the challenges and solutions for governing Artificial Intelligence at scale. The core objective was to demonstrate how specialized tooling can bridge the gap between rapid AI development and the growing demands for robust governance, risk management, and regulatory compliance.

The webinar highlighted that as organizations increasingly deploy hundreds or even thousands of AI models, traditional manual oversight becomes untenable. This creates significant risks, including compliance failures, reputational damage, and the silent failure of AI products. The speakers introduced a structured framework for AI governance, emphasizing the need for a dedicated governance layer that operates in tandem with the standard AI development lifecycle. This framework is built on three essential pillars: Lifecycle Management, Risk Management, and Compliance Management.

A key focus of the presentation was a detailed analysis of the current AI governance tool landscape, which was segmented into five distinct archetypes: Hyperscalers, Integrated AI Platforms, MLOps Tools, LLMOps Tools, and dedicated Governance Tools. To make the discussion tangible, the speakers presented two case studies: IBM's watsonx.governance, an enterprise-grade integrated platform, and Deeploy, a nimble MLOps-native solution. These examples illustrated how different tools cater to diverse organizational needs, from large corporations requiring comprehensive, multi-faceted solutions to smaller, agile teams focused on developer experience and specific compliance requirements like the EU AI Act.

The session concluded with a practical, five-step approach for selecting and implementing the right tooling, followed by an engaging Q&A that addressed critical questions about the selection process, stakeholder involvement, and managing regulatory complexity across different geographies. The overarching message was clear: embedding trust and governance into the foundation of AI development is not a barrier to innovation but a critical enabler for scaling AI faster, safer, and more effectively.

Nemko Digital

Overview and Objectives

This Nemko Digital webinar explored how organizations can scale AI responsibly by leveraging the evolving landscape of AI governance tools and technologies. The session aimed to:

- Explain why scaling AI is challenging from technical, organizational, and regulatory perspectives.
- Show how to move from a developer-centric lifecycle to organization-wide AI governance.
- Map the current tool landscape and share a practical selection framework.
- Deep-dive into two representative platforms: IBM watsonx.governance and Deeploy.
- Provide pragmatic guidance for selecting and adopting governance solutions at scale.

The Core Challenge: Why Scaling AI Responsibly is Difficult

The webinar began by outlining the complex landscape that makes scaling AI a significant challenge for modern organizations. Pepijn van der Laan explained that the difficulty arises from a confluence of factors that pull organizations in different directions.

First, the sheer speed of innovation in AI is unprecedented. The computational power (FLOPs) used for training cutting-edge models is exploding, leading to ever-more powerful and complex capabilities. This rapid advancement means that governance frameworks must be agile and forward-looking. Simultaneously, the regulatory landscape is evolving just as quickly. Governments worldwide are introducing new legislation, such as the EU AI Act, to mitigate the risks associated with powerful AI systems. This creates a moving target for compliance and requires constant vigilance.

Furthermore, the proliferation of AI means it is being embedded in everything from internal processes to consumer-facing products. This creates a massive inventory of AI systems that need to be tracked, monitored, and governed. The speakers highlighted several common pain points that keep executives up at night:

- *How can we convince our Chief Risk Officer that all AI risks are under control?*

- *How do we manage compliance when we have hundreds of AI-powered products?*
- *How can we ensure our AI products don't fail silently in production?*
- *How do we maintain control over products that use agentic AI with autonomous capabilities?*
- *How do we ensure our AI remains reliable and fair as the environment and data change over time?*

These questions underscore the central problem: manual approaches to AI governance are no longer viable. As Bas Overtoom noted, "When you go through a lot of AI models, just putting in the right guardrails and trying to monitor that in a manual way is becoming impossible." This scalability issue is the primary driver for adopting specialized AI governance tools.

A Framework for a Scalable Solution: The Governance Layer

To address these challenges, the speakers proposed a conceptual framework that separates the AI lifecycle into two distinct but interconnected flows: the Development Flow and the Governance Flow.

The Development Flow is the traditional, technically-focused process that data scientists and ML engineers are familiar with. It encompasses everything from use-case onboarding and data preparation to model training, deployment, monitoring, and eventual retirement. This flow is well-supported by a mature ecosystem of MLOps tools that streamline and automate the technical pipeline.

However, this flow alone is insufficient for ensuring responsible AI at scale. The webinar argued for the necessity of a Governance Flow, which sits above the development pipeline and addresses the broader organizational, risk, and compliance requirements. This layer is where business stakeholders, risk managers, legal teams, and auditors interact with the AI lifecycle. Key activities in this flow include:

- Use-Case and Risk Approval: Formal processes for evaluating the viability and potential risks of new AI initiatives before development begins.

- Risk Evaluation and Monitoring: Ongoing assessment of AI systems against predefined risk criteria and ethical principles.
- AI Inventory Management: Maintaining a centralized, up-to-date registry of all AI models across the organization.

This dual-flow model highlights the need for tooling that can connect these two worlds, providing a common platform for both technical and non-technical stakeholders. The benefits of such tooling fall into three main categories:

1. Lifecycle Management: Provides a holistic view of the entire AI inventory, enabling organizations to track model performance, detect data drift, and manage bias, not just within a single platform but across a potentially fragmented technology stack.
2. Risk Management: Creates a systematic way to track, mitigate, and report on AI-related risks. It offers visibility into vulnerabilities and allows for the implementation of standardized controls and policies.
3. Compliance Management: Automates the process of adhering to regulatory requirements. These tools often come with pre-built frameworks for major regulations (like the EU AI Act), streamlining documentation, audit trails, and reporting.

Navigating the AI Governance Tool Landscape

Pepijn van der Laan provided a detailed breakdown of the diverse and often confusing market for AI governance tools. He categorized the available solutions into five main archetypes, each with a different origin story and focus:

1. Hyperscalers: These are the major cloud providers (e.g., AWS, Azure, Google Cloud) who are typically fast-followers, adding governance features to their extensive suite of services to create a one-stop-shop experience.
2. Integrated AI Platforms: These platforms, like IBM's watsonx, offer an opinionated, end-to-end architecture that aims to balance a strong developer experience with comprehensive governance capabilities.
3. MLOps Tools: These tools originated to serve the needs of data scientists and ML engineers. Many, like the case-study example Deeploy, have evolved to incorporate broader governance features as they recognized its importance.

4. LLMOps Tools: A newer category focused specifically on the unique challenges of managing Large Language Models (LLMs), with a strong emphasis on developer experience and cost control.
5. Governance Tools: These are specialized solutions that focus primarily on the AI inventory, policy management, and compliance aspects, often with a close link to professional services.

To help organizations navigate this landscape, the speakers presented a seven-point selection framework for evaluating potential tools. This framework encourages a holistic assessment beyond just features and pricing:

Criteria	Key Considerations
Market Presence	Is the vendor established? Do they have a proven track record and strong backing?
Tool Capabilities	Does the tool meet functional requirements? What is the quality and ease of configuration?
Innovation	What is the product roadmap? Is the vendor's vision aligned with your future needs?
Support	What are the SLA commitments? What onboarding and training resources are available?
Integration	How well does it fit with your existing tech stack? Are there pre-built connectors?
Compliance	Does the tool support relevant local and industry-specific regulations?
Total Cost of Ownership (TCO)	What are the full costs (license, maintenance)? What is the risk of vendor lock-in?

This structured approach helps ensure that the selected tool not only meets immediate needs but also aligns with the organization's long-term strategy and existing infrastructure.

Case Studies in Contrast: IBM watsonx.governance and Deeploy

To bring the theoretical discussion to life, the webinar delved into two specific tools that represent different archetypes and approaches.

Case Study 1: IBM watsonx.governance - The Enterprise-Grade Platform

Representing the Integrated AI Platform archetype, IBM watsonx.governance is designed as a comprehensive, enterprise-wide solution. Backed by the global IT powerhouse, it aims to be a single toolkit to "direct, manage, and monitor your AI." Key characteristics include:

- Broad Integration: It is part of the wider watsonx AI & Data platform but can also integrate with third-party solutions from AWS, Microsoft, and others, supporting both cloud and on-premise deployments.
- Full Lifecycle Governance: The platform is built around a central AI inventory that provides a single source of truth for all models. It includes features for risk categorization, compliance management, policy enforcement, and automated workflows.
- Detailed Fact-Checking: A standout feature is the concept of model fact sheets. These documents automatically capture metadata and metrics throughout the model lifecycle, creating a detailed, auditable record for each AI asset. This is crucial for demonstrating compliance and explaining model behavior.

IBM watsonx.governance is positioned for large organizations that need a robust, scalable, and highly structured approach to AI governance, with a strong emphasis on auditability and control.

Nemko Digital

Case Study 2: Deeploy - The Agile MLOps-Native Solution

In contrast, Deeploy represents the MLOps Tool archetype that has expanded into governance. As a Dutch startup founded in 2020, Deeploy comes from a more developer-centric world. Its evolution reflects the market's growing demand for embedded governance. Key characteristics include:

- Developer-Friendly Origins: The platform started with a focus on responsible model deployment, performance monitoring, and explainable AI. This heritage is visible in its user interface and Python client, which are designed to fit seamlessly into data scientists' workflows.
- Compliance-Aware: Being based in the Netherlands, Deeploy has a strong focus on compliance with the EU AI Act. It has integrated features like audit trails, model tracking, and governance checklists directly into the deployment process.
- Flexible Deployment: It is available as a SaaS solution or for private cloud deployment on AWS and Azure marketplaces, and it integrates with popular data platforms like Databricks.

Deeploy is well-suited for organizations, particularly in regulated industries like finance and healthcare, that are looking for an agile, developer-friendly tool that embeds governance directly into the MLOps pipeline without creating excessive overhead.

A Practical Path Forward: The 5-Step Implementation Approach

Bas Overtoom concluded the main presentation by outlining a structured, five-phase approach that organizations can follow to select and implement the right AI governance tooling.

1. Explore: This initial phase is about understanding the landscape. It involves identifying ambitions, needs, and constraints, and collecting internal and external best practices to create a longlist of potential tools.
2. Assess: In this phase, the longlist is narrowed down to a shortlist. This involves prioritizing user stories and pain points, assessing the current architecture, and holding validation sessions with stakeholders to build a solid business case.
3. Select: This is the formal evaluation stage. It typically involves running a full RFI/RFP process, conducting a proof-of-concept (PoC) with 2-3 prioritized tools, and negotiating contract terms.
4. Activate: Once a tool is selected, this phase focuses on implementation. It includes system and process integration, tool configuration, and the crucial step of training and onboarding users.
5. Evolve: The journey doesn't end with activation. This final phase is about continuous improvement, updating the tool and processes as the organization's needs and the regulatory landscape evolve.

This methodical process ensures that the investment in AI governance tooling delivers real value and is successfully adopted across the organization.

Insights from the Q&A Session

The webinar concluded with a lively Q&A session that provided additional practical insights:

- On the Tool Selection Process: Pepijn emphasized the importance of keeping the selection process focused. Rather than a large, democratic committee, he recommended a small group of key stakeholders representing the full ecosystem (e.g., data science, risk, legal). He also stressed the importance of doing a thorough pre-selection to narrow down the options early, which contains the effort and duration of the process.
- On User Personas: When asked who the primary users of these tools are, Pepijn explained that they are designed to be a bridge between different communities. Data scientists and AI developers interact with the tool as part of their development and deployment workflows. In parallel, risk managers, auditors, and compliance teams use the same tool to check boxes, grant approvals, and

ensure that use cases are compliant. The tool thus serves as a common ground where these different worlds can collaborate.

- On Managing Regulatory Complexity: A key question addressed the challenge of managing different regulatory frameworks across various geographies and business units. Pepijn noted that this complexity is precisely why tooling is essential. Modern governance tools are increasingly configurable, allowing organizations to set up multiple risk and compliance frameworks (e.g., one for the EU AI Act, another for a different regional regulation). The tool can then orchestrate which framework applies to which AI solution, providing a centralized way to manage a complex and fragmented regulatory environment.

Conclusion: Building with Trust as the Foundation

The central message of the webinar was a powerful call to action: organizations must shift from "bolting on" trust as an afterthought to building with it as a foundational element of their AI strategy. The speakers convincingly argued that investing in a robust AI governance framework, supported by the right tooling, is not a cost center or a compliance burden. Instead, it is a critical enabler of innovation that allows organizations to scale their AI initiatives faster, more safely, and with greater confidence.

For attendees, the key takeaway is the need to move beyond ad-hoc, manual processes and adopt a systematic approach to AI governance. The provided frameworks for understanding the tool landscape and for selecting a solution offer a clear and practical roadmap for this journey.

Next Steps for Attendees

Nemko Digital offered several resources for those looking to continue their learning journey:

- Personalized Exploration: A 30-minute one-on-one call to evaluate specific requirements and explore how Nemko Digital can provide support.

Nemko Digital

- LinkedIn Community: An invitation to join the Nemko Digital LinkedIn community to stay informed about upcoming webinars on topics like the EU Data Act, AI Trust Marks, and the synergies between cybersecurity and AI.
- Further Information: Attendees were encouraged to visit the Nemko Digital website for more details on their services.

By embracing the principles and practices outlined in this webinar, organizations can transform AI trust from a challenge to a powerful competitive advantage.