SPEEDNET

Tutorial

I have implemented an AI system. What's next?

9 AI governance elements you cannot overlook

Table of contents

Executive Summary 3

Introduction 4

The new reality after AI implementation 5

9 key pillars of mature AI governance 6

Visualisation of the 9-pillar model 6

Stage I: Building foundations and taking control 7

Pillar 1: Establishing the organisational framework and accountability 8

Pillar 2: Implementing data governance 10

Pillar 3: Developing organisational culture and ethical awareness 12

Pillar 4: Implementing technological oversight 14

Stage II: Implementing strategy and managing risk 16

Pillar 5: Implementing proactive risk management 17

Pillar 6: Implementing independent verification and audit 19

Pillar 7: Ensuring compliance with the legal framework 21

Stage III: Optimising and verifying 23

Pillar 8: Measuring effectiveness and optimising oversight 24

Pillar 9: Verifying strategy and return on investment (ROI) 26

Conclusions and strategic recommendations 28

The next step: The path to automated governance 29

References and further reading 30

About Speednet 31

Executive summary

The rapid adoption of artificial intelligence reveals a deep conflict: despite huge investments, the average return on AI implementations remains at a surprisingly low 5.9%.

Numerous market examples show that the main reason for this situation is ineffective governance. In most cases, the problem is not the technology but a lack of organisational readiness. Companies invest in algorithms while neglecting the processes required to ensure their effectiveness, security, and legal compliance.

These risks translate directly into specific costs: uncontrolled use of AI by employees increases the average cost of a security breach by $670,000, while the annual manual cost of maintaining compliance for a single high-risk system under the EU AI Act is between €116,000 and €137,000. As a result, many AI initiatives, rather than generating profit, become a hidden source of expense that erodes the initial investment's value. This paralysis clearly demonstrates that traditional, reactive risk management models are insufficient for AI. Organisations worldwide clearly need a new, planned approach to balance innovation with security.

In response to these urgent challenges, we present a specific action plan: a proven 9-pillar AI governance model. We demonstrate, step by step, how to take control of artificial intelligence within the company, starting from establishing order in the chaos and moving towards achieving full optimisation.

We also prove that AI governance does not have to be merely a costly obligation; it can become a tool for faster and safer innovation.

Introduction

The AI investment paradox. Why are huge expenditures not translating into profits?

Implementing artificial intelligence is only the beginning. The real challenge lies in effectively managing this technology, which requires ensuring appropriate oversight and operating in compliance with new regulations.

This report shows how companies build a real, sustainable competitive advantage through mature AI governance.

Although AI adoption is progressing rapidly (with **78% of organisations** already having implemented the technology according to Stanford HAI), market data reveals a fundamental contradiction.

As IBM research indicates, the average return on investment for AI initiatives is just **5.9%**. This low return, despite high adoption, reveals a critical gap between the pace of implementation and the level of organisational preparedness.

This gap in governance translates directly into real financial losses. In fact, negligence in managing artificial intelligence carries many threats, including financial ones. This is confirmed by data from the IBM report, which states that uncontrolled use of AI tools by employees results in an increase in the average cost of a security breach by **$670,000**.

At the same time, new regulations, such as the EU AI Act, are introducing severe financial penalties.

Analysis by Holistic AI shows that these can reach **as high as €35 million or 7% of a company's global turnover.** Beyond penalties, hidden costs related to inefficiency can pose an equally significant problem.

According to Speednet analysis, the average annual cost of maintaining compliance for a single high-risk AI system in the EU banking sector ranges **from €116,000 to €137,000**.

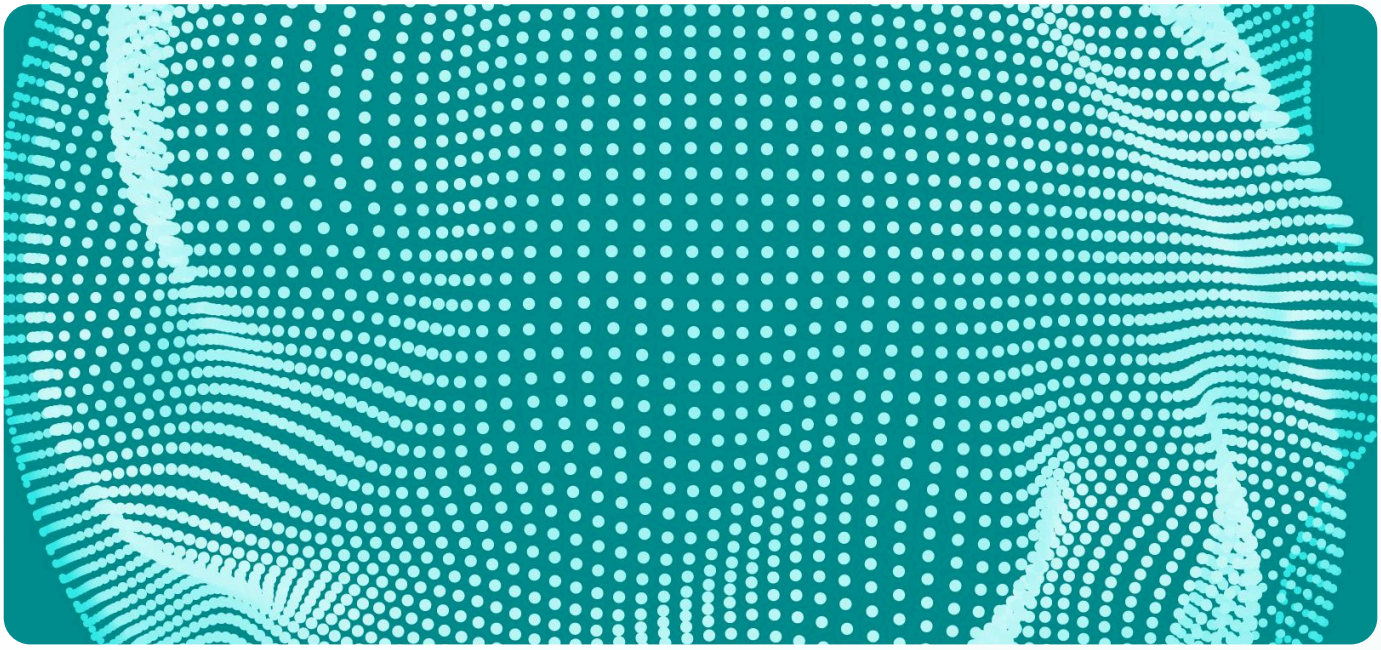
The largest portion of this amount, about **36%**, is allocated by companies to testing and validating models.

These figures, multiplied by the number of systems in an organisation, clearly show that a manual approach to governance is not only risky but also financially unfeasible.

In this document, we present 9 key pillars of mature AI systems governance, grouped into t**hree logical stages:** from building foundations, through implementing strategy and mitigating risk, to optimisation and verification.

STAGE I	Building foundations and taking control
STAGE II	Implementing strategy and managing risk
STAGE III	Optimising and verifying

The new reality after AI implementation



Artificial intelligence, as defined by the IIA AI Auditing Framework, involves systems equipped with intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalise, or learn from experience.

The current boom in AI applications, especially advanced language systems, is revolutionising how businesses operate.

"The 2025 AI Index" report by Stanford HAI forecasts that the global AI market will grow fivefold within the next five years, and the technology's adaptation in business has become widespread, already encompassing 78% of companies.

This is precisely why mature governance of artificial intelligence is essential to protect and build company value. This principle is confirmed by analyses from McKinsey.

Many companies make the mistake of trying to simply "attach" an AI module to an existing, often inefficient, process, expecting an automatic improvement. This approach overlooks a key fact: real benefits come not from automation itself, but from the fundamental change in working methods that this technology enables. Without this deep restructuring, artificial intelligence becomes merely a costly and complicated addition to the old problem. It does not solve it at the source and ultimately fails to deliver the expected profits.

9 key pillars of mature AI governance

An effective artificial intelligence governance system rests on nine key pillars.

This is not a random collection of tasks, but a clear, well-planned roadmap.

We have divided it into three logical stages that, step by step, lead an organisation from chaos to full control and optimisation.

Understanding two principles here is essential: chronology and completeness. The individual stages build upon each other: one cannot effectively optimise processes (Stage III) without first building solid foundations (Stage I) and giving actions a strategic direction (Stage II).

Similarly, each of the nine pillars is an essential part of the whole. Ignoring even one is like building a bridge with a missing key span; the entire structure becomes unstable and risks collapse.

Only by implementing all pillars comprehensively can we guarantee the creation of a resilient and effective AI governance system.

STAGE I

Building foundations and taking control

- Pillar 1 Establishing the organisational framework and accountability
- Pillar 2 Implementing data governance
- Pillar 3 Developing organisational culture and ethical awareness
- Pillar 4 Implementing technological oversight

STAGE II

Implementing strategy and managing risk

- Pillar 5 Implementing proactive risk management
- Pillar 6 Implementing independent verification and audit
- Pillar 7 Ensuring compliance with the legal framework

STAGE III

Optimising and verifying

- Pillar 8 Measuring effectiveness and optimising oversight
- Pillar 9 Verifying strategy and return on investment (ROI)

STAGE I STAGE II STAGE III

Building foundations and taking control

At this stage, we concentrate on gaining control over what is already operating within the organisation. Our goal is to build solid organisational, procedural, and ethical foundations. Without this basic groundwork, our future efforts may yield very limited results.

Pillar 1	Establishing the organisational framework and accountability
Pillar 2	Implementing data governance
Pillar 3	Developing organisational culture and ethical awareness
Pillar 4	Implementing technological oversight

Pillar **1**

Establishing the organisational framework and accountability

Why is this important?



The absence of a formal framework is a direct cause of real losses. Data from the IBM report is clear on this: 63% of organisations that experienced a security breach in their AI systems had no formal AI governance policy in place. Furthermore, almost all (97%) admitted they lacked appropriate access control mechanisms.

This is why establishing a clear structure, transparent rules, and a distinct division of responsibilities is the absolute foundation.

What are the key steps?



We implement a governance structure.

We base it on a proven standard, such as the Three Lines Model (in accordance with the IIA AI Auditing Framework). In practice, this means the **first line** (business and development teams) directly manages the risk of their AI systems in daily operations. The **second line** (risk and compliance departments) creates company-wide frameworks and oversees compliance, acting as an independent controller. The **third line** (internal audit) provides the board with an objective, independent assessment of whether the entire system functions effectively.

We create a central inventory of AI systems.

We conduct an internal audit to map and catalogue all AI applications in the company. This inventory must be a dynamic tool, not a static register, defining each system's business owner, the key datasets it operates on, its formal risk level assessment (e.g., according to EU AI Act categorisation), its purpose, and its key metrics. Without such a central inventory, we do not know what is actually under control, and governance becomes reactive and ineffective.

We develop an acceptable use policy.

This document defines and communicates clear "rules of the game". The policy must precisely specify which categories of AI tools (especially external ones, like generative models) are permitted and which are prohibited, how employees must handle confidential and company data when interacting with AI, and what ethical boundaries must not be crossed. The goal is to ensure innovation does not come at the expense of security.

We designate accountability.

We create a clear document, often a RACI matrix, that defines the roles of individuals in key processes. This ensures everyone knows who is Responsible for daily model monitoring, who is Accountable for the final decision in an incident, who must be Consulted before a change, and who merely needs to be Informed of its results. This effectively eliminates the problem of 'unclear' responsibility.

We appoint a cross-functional AI team.

Finally, to ensure coordinated action across the organisation, we form an AI Steering Committee. It must include representatives from key areas: technology (CIO/CTO), legal, risk management, key business units, and HR (due to AI's impact on employees). Such a diverse team can effectively assess AI's impact from every perspective while overseeing strategy implementation.

What are the risks of inaction?

Ignoring this pillar exposes the company to a high probability of security breaches, uncontrolled risk from the spread of unapproved AI tools (which IBM reports increases breach costs by an average **of $670,000**), and accountability gaps that make it impossible to identify the responsible party during an incident.

Expert advice

Attempting to manage all AI initiatives through a single, central decision-making committee inevitably leads to innovation paralysis.

The most aware organisations are moving away from this model towards a structure that smartly divides responsibility into two complementary levels: strategic and operational. At the strategic level, a specialised team of experts operates, usually as a Center of Excellence (AI CoE).

This team does not get bogged down in daily project tasks. Its role is to create solid foundations for the entire organisation: it sets technological standards, defines security policies, and determines acceptable risk boundaries.

It acts as an architect providing precise plans, not a construction manager supervising every move. The full burden of project execution rests with small, dynamic teams integrated into business departments like marketing, finance, or operations. Delegating authority this way dramatically accelerates progress, but it also requires significant organizational maturity. Precisely defining the autonomy of these teams becomes essential to avoid competency chaos and friction with the central team.

Pillar **2**

Implementing data governance

Why is this important?



With a framework for people and accountability in place (Pillar 1), the next logical foundation is governing the asset they use most: data. AI algorithms act as an amplifier: every error or historical bias in the input data will not only be replicated but also amplified and scaled. Therefore, feeding low-quality data into systems leads directly to creating flawed tools. For example, if historical credit data contained biases, an AI system will not just repeat them; it will make them the foundation of its decision-making logic and begin to systematically discriminate against entire customer groups on a massive scale. Organising this area is not just a technical issue; it is the foundation for a profitable and ethical business.

What are the key steps?



We ensure data quality and consistency.

We implement processes for ongoing checks of information feeding the models, including automated completeness and accuracy tests, which allows for correcting errors at the source before they can influence the algorithm's "decisions".

We establish control over the full data lifecycle.

We define clear rules regarding how data is collected, processed, stored, and deleted, ensuring full compliance with regulations, especially GDPR.

We implement strict data access control.

We apply the principle of least privilege, guaranteeing that only authorised individuals and systems have access to specific datasets. As the IIA AI Auditing Framework highlights, implementing mechanisms for the segregation of duties is also critical here to prevent potential misuse.

We pay special attention to privacy protection.

We actively use techniques such as anonymisation and pseudonymisation to effectively protect personal data and other confidential information, especially in systems accessible to a wide range of users.

We document and track data lineage.

To ensure full transparency, we implement solutions that allow us to gain insight into the data based on which the system made a specific decision. Full data lineage is essential for conducting audits and effectively building trust.

What are the risks of inaction?



A lack of data governance leads to a series of problems: from financial losses resulting from flawed business decisions, through the significant reputational risk of algorithms replicating historical biases, to actual financial penalties for GDPR violations. What is worse, these problems often build up quietly. An AI model "fed" bad data does not break down overnight. It simply begins to make subtly incorrect decisions, and initially, no one notices. By the time the company realises its sales forecasts or customer recommendations have long been flawed, it has usually already faced months of losses that are difficult to recover.

Expert advice

Mature organisations do not react to data problems; they prevent them using data contracts. These are automated, code-based agreements defining the required standard at the interface between systems. If source data (for example, from a CRM system) violates these rules, the contract is broken, and the process is automatically blocked. This protects the AI model from being "poisoned" by faulty information.

This is a key shift that moves responsibility for data quality from analytical teams back to the data source. It forces discipline at the foundation and guarantees that only data compliant with the standard reaches analysis.

This approach is fundamental to modern thinking about data as finished, reliable products. Instead of manually verifying the quality of input data received from other departments each time, analytical teams receive a verified resource with a quality guarantee.

This shortens the time needed to build and deploy new AI models, as it eliminates the most time-consuming and repetitive phase of work: finding, preparing, and cleaning data.

Pillar **3**

Developing organisational culture and ethical awareness

Why is this important?



A solid framework (Pillar 1) and clean data (Pillar 2) are not enough if the people using the technology are not prepared. Ultimately, the effectiveness of even the best technologies and procedures always depends on the human factor. It is the attitudes, understanding, and trust of employees that determine the success of the entire initiative. An ideal governance framework will remain useless if the team does not understand the risks or begins to resist change. Actively developing awareness and promoting responsible attitudes is a key element of the entire system.

What are the key steps?



We define and implement AI ethical principles.

Our principles must be specific and measurable. We therefore base them on the characteristics that the IIA AI Auditing Framework identifies as key for any system: explainability, fairness (lack of bias), reliability, accountability, and transparency. This is how we create clear standards to evaluate our technologies.

We implement training programmes for all employees.

We organise mandatory workshops focusing on two areas: the practical use of tools and, more importantly, understanding and applying the company's AI policy. We explain how to recognise common risks (e.g., related to data privacy) and present clear procedures for reporting any incidents or concerns.

We communicate actively and manage the change efficiently.

We conduct open communication to build understanding of all new processes. Leaders at every level must actively explain the goals and principles of AI governance to employees. In this way, we help our staff adapt to something that is unknown and may cause concern.

We create secure channels for reporting concerns.

Creating procedures is the first step. The second, more important one, is encouraging people to use them. To this end, we launch simple, anonymous communication channels. They allow employees to report their doubts or incidents without fear of negative consequences.

What are the risks of inaction?



A lack of attention to culture and awareness leads to a series of linked problems. It all starts with human error: an untrained employee may misinterpret AI results and make a mistake, leading to an incident. This, in turn, causes the company to lose the trust of customers and employees.

This loss of trust is not a theoretical image problem. As numerous market examples show, it translates directly into financial results. Customers who feel treated unfairly by an algorithm may leave for a competitor. At the same time, employees who are told to rely on tools they simply do not understand will quickly find ways to bypass them. This will completely wipe out the intended return on investment.

Expert advice

Traditional AI ethics training fails because the theory is detached from practice. This is why market leaders adopt a two-pronged approach to building responsible technology.

On one hand, they appoint specialised, internal ethical red teams, which, like "ethical hackers," test company systems, searching for hidden biases and vulnerabilities.

On the other hand, they involve the entire organisation by extending existing bug bounty programmes to include rewards for detecting ethical problems. This combination of specialist auditing and employee engagement means AI ethics ceases to be an abstract concept and becomes a real, measurable element of the company culture.

This model has one more key advantage: it completely changes the company's position in discussions with customers and regulators. Instead of just making declarations and presenting training certificates, the organisation can prove it has an active, internal mechanism for detecting and correcting its own errors. Proactively searching for ethical problems before they become public incidents builds a high level of trust that cannot be faked. It is the strongest proof of taking responsibility seriously.

Pillar **4**

Implementing technological oversight

Why is this important?



A strong culture and clear principles (Pillar 3) are essential, but they cannot be enforced manually. Rules and procedures alone, while often looking good on paper, are not enough without the technology to back them up. Without the right tools to enforce them, they provide only a false sense of security. This pillar, therefore, builds the necessary technological backbone that allows us to move from reacting to problems to proactive, automated oversight. It is technology that, in practice, transforms our strategy into a functioning and measurable system.

What are the key steps?



We provide appropriate IT resources, which is essential for professional activities in the AI domain.

This includes not only the infrastructure itself (computing power for training models and storage for data) but also, a frequently overlooked point, the highly qualified specialists capable of managing it. Investing in hardware without investing in people is a direct path to frustration and wasted money.

We implement process automation tools (MLOps).

These platforms allow us to standardise and automate the entire model lifecycle, from creation to deployment. In practice, MLOps significantly shortens the time needed to deploy a new model, giving teams more time for innovation rather than struggling with the deployment process.

We launch continuous monitoring systems.

Once deployed, an AI model is a dynamic system; its effectiveness can decline over time. Continuous monitoring allows us to automatically track key performance indicators and receive alerts before the model starts making decisions that mislead our employees or customers.

We implement automated data lineage tools.

In the event of a regulatory audit, instead of a nervous, weeks-long search for evidence, these systems allow us to instantly prove what data we trained a given model on.

We supplement our toolkit with explainable AI (XAI) solutions.

An AI's decision alone is often not enough, as people naturally distrust processes they do not understand. Both customers and internal managers must understand *why* the system made a specific decision, which is essential for building trust.

What are the risks of inaction?



A lack of technological foundations leads directly to serious operational disruptions. Initially, this shows up as low efficiency, slowing down key processes in the company; manual model management quickly begins to block any development. The worst part, however, is what remains unseen. We are talking about the slow decline in algorithm accuracy, which means we might be making flawed business decisions for months without even knowing it.

Expert advice

Efficient AI implementation requires more than data access; it requires certainty that the technology operates within established rules. This guarantee is provided by an approach known as governance as code, which transforms company standards and policies into a set of active, automated rules. Thanks to this, guidelines (for example, concerning model bias limits or documentation requirements) cease to be static text. The system automatically ensures compliance at every stage of work, taking this burden off analysts and developers.

In this way, the organisation solves the fundamental conflict between speed and control. The compliance department stops acting as a "brake," manually checking projects just before deployment. Instead, it becomes a partner providing development teams with control mechanisms. This gives analysts freedom and autonomy, allowing them to experiment quickly in a secure, automatically controlled environment.

STAGE I STAGE II STAGE III

Implementing strategy and managing risk

Now that we have control over the foundations from Stage I (people, data, culture, and technology), we can move from establishing order to strategic actions.

At this stage, we focus on consciously shaping AI's role in our organisation.

The goal is to ensure our artificial intelligence initiatives are closely aligned with business objectives, that we proactively manage their associated risks, and that we operate in full compliance with the legal framework.

Pillar 5	Implementing proactive risk management
Pillar 6	Implementing independent verification and audit
Pillar 7	Ensuring compliance with the legal framework

Pillar **5**

Implementing proactive risk management

Why is this important?



With our foundations in place, we can now address risk, not reactively, but proactively. Many companies approach AI-related risk reactively, analysing problems only when they escalate to the point of causing real losses. This is a costly and dangerous strategy, especially since AI system errors often magnify quietly and over long periods. This pillar completely changes that approach. We do not wait for problems to occur. Instead, we systematically identify, assess, and mitigate them before they can harm our operations.

What are the key steps?



AI-related risks must become an integral part of the company-wide risk management system.

To this end, we enter them into the central risk register, establish an acceptable risk level for AI initiatives, and assign responsibility for them to business managers.

We analyse high-profile cases and other companies' errors.

These are a valuable source of knowledge. We study known problems, such as discrimination in recruitment algorithms (Amazon) or in credit limit allocation (Apple Card). The lessons learned serve as a practical checklist for our own safeguards.

We conduct a systematic AI impact assessment for every new AI system before it is deployed.

Such analysis covers the impact on data privacy, potential for misuse, fairness towards different social groups, and resilience to attacks. This allows us to plan and implement remedial actions before the system goes live.

We manage the risk associated with a lack of transparency in models whose operation is difficult to interpret.

We cannot simply accept this fact. Our task is to actively identify these gaps and then, following IIA AI Auditing Framework recommendations, introduce additional safeguards, for example, by implementing human oversight for the most critical decisions.

We conduct continuous review and updating of risk maps.

Our risk assessment must be an ongoing process, as the threat landscape is constantly changing. This requires us to react to new regulations, high-profile incidents in our industry, or the planned use of AI in new, sensitive areas (like HR).

What are the risks of inaction?



Initially, the greatest threat is "invisible" risks. These are subtle errors. They build up quietly, often over many months, leading to cumulative losses. We usually discover them when it is already too late. When the problem finally becomes visible, the company is completely unprepared for it. A chaotic, ad-hoc reaction often only worsens the situation. In a legal dispute, a company with no documented risk assessment will find it difficult to prove it took appropriate steps to prevent problems.

Expert advice

To give risk assessment a real dimension, one must move from general descriptions to hard financial data.

Instead of using a qualitative "high/medium/low" scale, experts implement "probabilistic financial risk modelling".

They use statistical methods, such as Monte Carlo simulations, to estimate the expected annual loss (in currency) for a specific AI system failure.

The question changes from "is this risky?" to "what is the probable revenue loss if our recommendation engine's accuracy drops by 3% over the year?".

This financial model allows abstract risk to be translated into the specific language of finance, which in turn makes it easier to prioritise corrective actions and precisely calculating the return on investment from mitigating it.

Pillar **6**

Implementing independent verification and audit

Why is this important?



Proactively managing risk (as in Pillar 5) is the first half of the solution. The second half is proving that those management controls actually work. Guarantees do not come from simply implementing control mechanisms. We must also gain certainty that these implemented controls do not lose their effectiveness over time. Independent verification is the next critical step, allowing us to objectively confirm the effectiveness of our actions.

What are the key steps?



We define the tasks of internal audit.

Internal audit has two primary tasks. First, it acts in an advisory capacity, helping to design effective control mechanisms from the very beginning. Second, it acts in a verification capacity, conducting formal audits that assess whether the implemented controls are actually working.

We develop a regular AI audit plan.

This plan is based on risk analysis. This means we first concentrate on the AI systems with the greatest potential impact on the company's operations. The plan is regularly updated to account for new implementations and changing threats.

We conduct in-depth verification and testing.

During an audit, we go a step further than just checking documentation. Auditors must verify whether the input data feeding the models is complete and accurate. They also test whether the results generated by AI are consistent with the business objectives for which the system was created.

We present results directly to the board and the audit committee.

We present concise and understandable reports to top management. Each report must clearly indicate identified weaknesses and include specific recommendations. We assign a responsible person and a completion deadline to each recommendation.

What are the risks of inaction?



Ignoring independent verification leads to serious, escalating threats. The greatest of these is a false sense of security. We may be convinced our control mechanisms are working when, in reality, they have long been ineffective. This, in turn, leads to a loss of credibility in the eyes of partners and regulators. We are then unable to objectively prove that we take AI governance seriously. The problem runs deeper, however. A lack of external verification inevitably leads to a decline in internal discipline. Development teams, working without systematic auditing, begin (often in good faith) to take shortcuts, such as skipping edge-case testing. An audit thus serves not only to find errors; its very presence acts as a deterrent and is essential for maintaining high standards of quality and ethics throughout the organisation.

Expert advice

For an audit to be specific and technical, not merely procedural, it is worth implementing the concept of an "AI bill of materials" (AIBOM).

In practice, this means maintaining a detailed, versioned document for each key AI system, which functions as a detailed technical specification.

Such an AIBOM contains a list of all system "components": the exact model version, the datasets used for its training (along with their digital signatures), key training parameters, a list of open-source libraries used, and documented results of bias and performance tests.

In this model, the auditor's task is no longer just assessing the process on paper, but technically verifying the integrity and reproducibility of this bill of materials.

Using an AIBOM makes the audit fully measurable and significantly more reliable.

The value of an AIBOM, however, extends beyond mere reproducibility. It becomes a primary tool for managing the risk associated with the components that make up the AI system. Modern systems are rarely built from scratch; companies often use external datasets, pre-trained foundation models (e.g., open-source), or dozens of libraries. An AIBOM allows the auditor to precisely trace this "origin" and assess the inherited risk (for example, hidden biases in a supplier's data, security vulnerabilities in a specific library, or a model's licensing restrictions). This gives the company full insight not only into what it has built itself, but also into what it has downloaded from external sources and incorporated into its systems.

Pillar **7**

Ensuring compliance with the legal framework

Why is this important?



This internal and technical verification (Pillar 6) is critical, but it must operate within a rapidly changing external legal landscape. The EU AI Act introduces a new legal reality. Non-compliance carries multi-million euro penalties. This makes legal compliance a strategic matter, discussed at the board level, and a key element in protecting the company from losses. This pillar is a practical guide to this new reality.

What are the key steps?



We analyse gaps in our current systems.

We carefully compare our documentation and processes against the specific requirements of the AI Act. This is often tedious but necessary work. Its result is a precise task list, so we know exactly what we need to improve.

We map our AI systems, assigning them to risk categories.

This step is essential as it avoids unnecessary work. We must remember that the most strict requirements apply to "high-risk" systems, so this classification allows us to direct our resources where they are truly needed.

We apply the ISO/IEC 42001 standard as a practical implementation framework.

Instead of inventing processes from scratch, we rely on a proven standard. This standard provides us with a ready-made structure and specific guidelines. It also gives our teams confidence that they are building an AI management system in a manner compliant with legal requirements.

We automate the technical documentation creation process.

Manual documentation is inefficient and usually leads to costly errors, especially with a large number of models. We therefore implement tools that do this automatically, freeing up our teams.

What are the risks of inaction?

Problems with AI regulation compliance primarily risk severe financial penalties. Holistic AI analyses indicate these can reach as high as 7% of global annual turnover. That is not all. Regulators can also order the immediate shutdown of a system, which in practice means halting key business processes.



We often underestimate what such a "system shutdown order" means in practice. It is not just a temporary technical problem, but the threat of an immediate shutdown of core operations. Imagine a high-risk system that assesses creditworthiness being shut down overnight. The bank loses its ability to issue loans, striking at the very heart of its business. Moreover, such a public failure in the eyes of the regulator destroys the trust of investors and business partners, who will begin to see the company as an unstable partner. Rebuilding that reputation is usually many times more expensive than the financial penalty itself.

Expert advice

Chasing ever-changing regulations is a losing battle.

This is why intelligent companies reverse the logic: instead of creating more procedures, they implement a compliance by design approach.

This involves building a library of ready-made technological components, each with compliance and legal approval "built-in".

In practice, instead of ten teams interpreting GDPR on their own, the company provides them with one, ready-to-use anonymisation module.

Similarly, in the context of the AI Act, teams use a template that automatically completes most of the documentation.

Legal compliance is no longer a problem to remember; it becomes a built-in, automatic feature of the tools teams use every day.

STAGE I STAGE II STAGE III

Optimising and verifying

With a strategy in place and legal/risk processes under control (Stage II), the final stage allows us to achieve operational excellence.

The goal here is to measure the effectiveness of the implemented processes, their continuous optimisation, and independent verification of whether AI investments are delivering the expected return.

This allows our organisation to scale AI initiatives in a repeatable and secure manner, while adhering to all procedures.

Pillar 8	Measuring effectiveness and optimising oversight
Pillar 9	Verifying strategy and return on investment (ROI)

Pillar **8**

Measuring effectiveness and optimising oversight

Why is this important?



Implementing AI governance is a significant investment. The business will rightly ask: "what is in it for us?". If we cannot sufficiently justify our proposal, the business will always see our actions as an unnecessary cost. This pillar shows how to use simple data to prove the real value of governance and continuously improve it.

What are the key steps?



We define key performance indicators (KPIs).

We translate abstract concepts like "ethics" or "security" into specific, measurable indicators. These can be simple numbers, such as the percentage of models with full documentation, average incident response time, or the bias detection rate, as recommended by Relyance AI. Such a scorecard allows for objective progress tracking.

We measure return on investment (ROI).

We calculate it in two ways. On one hand, we show how much money we save by avoiding fines and legal problems. On the other, we prove how good governance helps bring new products to market faster and gain customer trust, which, as analyses from the California Management Review show, is a key element of value generation.

We translate technical data into business value.

We translate technical data into the language of finance. If monitoring detects a 5% drop in model accuracy, our job is to calculate the potential revenue loss this generates daily. We also analyse the total cost of governance, including specialists' time.

We identify bottlenecks.

The collected data allows us to find where we are wasting time and money. If our top specialists spend 40% of their time on repetitive checks, it is a signal that this area requires urgent automation.

What are the risks of inaction?



First, we generate hidden costs by wasting highly-paid specialists' time on repetitive, manual control tasks. Second, such a manual approach is unscalable and simply cannot be maintained across a growing portfolio of AI systems. Worst of all, without measurable data, we lose the chance for continuous improvement. Not knowing what works and what does not, we are unable to systematically improve our approach to governance.

Expert advice

The ultimate goal of corporate governance is not merely risk reduction, but primarily the acceleration of safe innovation deployment.

Therefore, the key metric to measure is the "velocity of innovation within governance".

Instead of focusing solely on the declining number of incidents, the organisation should track how the average time from an AI model idea to its production deployment (having successfully passed all automated control gates) is shortening.

A systematically decreasing time is the best proof that well-designed governance is becoming an accelerator, not a brake, for the business. This, in turn, provides the strongest argument for further investment in its development.

Pillar **9**

Verifying strategy and return on investment (ROI)

Why is this important?



Pillar 8 focused on the ROI of the governance process itself. This final pillar looks at the ROI of the AI initiatives that governance supports. Without this check, AI initiatives, when unlinked from business goals, can easily become a costly "research laboratory". They generate interesting technologies but bring no real value to the company. The goal of this pillar is to ensure our AI investments are not "art for art's sake". Instead, we ensure they translate into specific, measurable benefits for the organisation.

What are the key steps?



We involve senior management in defining goals.

We ensure that the goals for each AI initiative are defined in a Specific, Measurable, Achievable, Relevant, and Time-bound (SMART) way. For example, instead of the general goal "we will improve customer service," we define a SMART goal: "we will reduce the average response time to queries by 30% within 6 months by implementing an AI chatbot".

We develop an AI strategy consistent with company goals.

We create a document that clearly answers the question: how will specific AI actions support the achievement of the organisation's main business objectives? This strategy also defines the acceptable risk level for different types of initiatives, helping us avoid the risk of irresponsible or misguided decisions by individual teams.

We measure the tangible and intangible benefits of implementation.

We analyse not only the direct, easily calculated financial benefits (e.g., cost reduction, revenue growth, as confirmed by IBM data). We look just as carefully at intangible benefits, such as increased customer satisfaction (measured, e.g., by NPS) or improved employee satisfaction, as AI allows them to focus on more creative tasks.

We compare our results with the market.

We regularly analyse publicly available information on how other companies in our industry use AI and the value they achieve. This gives us an important reference point and allows us to assess whether our own results are satisfactory or if we are lagging behind the competition.

We aim to redesign processes to maximise value.

As McKinsey indicates, we achieve the greatest benefits from AI not when we automate an old process, but when we use the technology to completely rethink and redesign the way we work. This redesign is the ultimate goal, as it drives true transformation, not just cosmetic improvements.

What are the risks of inaction?



We see three main problems here. The first and most obvious is ineffective investment: we finance projects that never deliver a return. This, in turn, leads to a loss of support from the board, which, seeing no tangible benefits, becomes reluctant to fund further initiatives. As a result, the company loses its strategic direction, developing AI for technology's sake, with no connection to real business needs.

Expert advice

Many companies fall into the trap of financing dozens of small, isolated AI projects.

Each may look promising, but their combined impact on the company's financial results is often insignificant.

This is why more mature organisations are moving away from this model towards strategic, domain-based transformation.

This approach is based on a simple principle: instead of dividing the budget, the company consciously concentrates almost all its AI investment on the deep restructuring of a few (2-3) key areas, like marketing or logistics.

Each functions like a separate business unit with its own P&L and a precisely defined ROI target.

The business benefits flowing from such a deep transformation of one area are far greater than the sum of profits from dozens of minor improvements that get lost in the scale of the entire organisation.

Conclusions and strategic recommendations

The analysis of the nine pillars proves that organisations approach AI corporate governance in two ways.

Most treat it as a set of regulatory obligations and a cost to be incurred to avoid problems.

Market leaders, however, perceive it differently: they see it as a mechanism for building sustainable advantage.

The path we have presented shows how to move from the first group to the second.

During the **first stage (Foundations and control)**, we establish order from the chaos. We gain control over tools, data, processes, and people, creating a stable and predictable basis for further action.

At this stage, the goal is not yet innovation, but building a secure environment in which innovation can develop in the future.

Only on this basis, during the **second stage (Strategy and risk management)**, can we give our actions a strategic direction.

This is where we ensure our technology is secure, compliant with the law, and working towards real business goals, rather than being just a technological experiment.

Corporate governance begins to act as a compass, pointing in the right direction.

The **final stage (Optimisation and verification)** is dedicated to transforming corporate governance into a measurable and efficient system.

Here we prove its financial value and automate it to operate more cheaply and effectively.

Our goal is a situation where solid governance allows us to implement AI-based innovations faster and more confidently than our competitors.

The next step

The path to automated governance

Why is a manual approach not financially viable?

The manual implementation and maintenance of all nine pillars are extremely expensive, slow, and impossible to scale across the organisation. As Speednet* data shows, the average annual cost of maintaining compliance for a single high-risk AI system **ranges from €116,000 to €137,000**. This amount, multiplied by the number of models, clearly shows that the key to success in AI governance is not hiring more analysts, but smarter automation.

It is worth remembering that a manual approach is not viable because it creates a bottleneck and interrupts the automated MLOps cycle. Instead of a smooth deployment, our ML Engineers must halt work and manually perform validation (for example, generating reports on concept drift, checking the model's resilience to adversarial attacks, or documenting the origin of training data). As a result, we waste the time of highly qualified specialists on tasks that are simultaneously costly and prone to error.

How to justify the investment?

Do you need solid arguments to justify investing in an AI governance platform? Base your plan on concrete data.

Use our analysis, which will allow you to estimate potential savings and return on investment (ROI) based on your organisation's parameters.

Check our meta-analysis - a detailed calculations based on benchmarks from over 50 European firms in the financial sector.

Meta-analysis →

Want to move from theory to action?

Let's discuss how we can support your organisation in implementing these pillars.

Let's talk →

References and further reading

- THE IIA'S Artificial Intelligence Auditing Framework
- IBM: How to maximize ROI on AI in 2025
- IBM Report: 13% Of Organizations Reported Breaches Of AI Models Or Applications, 97% Of Which Reported Lacking Proper AI Access Controls
- Penalties of the EU AI Act: The High Cost of Non-Compliance
- The 2025 AI Index Report
- The state of AI in 2025: Agents, innovation, and transformation
- Shaping Europe's digital future: AI Act
- ISO/IEC 42001: a new standard for AI governance
- On the ROI of AI Ethics and Governance Investments: From Loss Aversion to Value Generation
- Measuring AI governance metrics that matter
- AI Governance Costs in EU Banking: Meta-Analysis

Authors



Tymoteusz Olszewski
Head of AI



Adam Rasiewicz
Senior Business Development Manager

SPEEDNET

Technology that drives modern finance



A shorter time to market and a more predictable ROI for your digital financial products, delivered through our outcome-based model and extensive BFSI sector experience.

mBank · VELO BANK · Santander · Bank Pekao · ALIOR BANK · Bank BPH · POP Pankki · GET IN BANK · ICA Banken · saar LB · Genossenschaftsverband Verband der Regionen · BANCOVO. · ringier axel springer · GENERALI · creditstar · paytree · cash·· · FICO · Wolters Kluwer

Our offering | Solutions | Delivery acceleration services

Solutions



Ready-made digital solutions designed for fast market impact

Our core portfolio brings together proven IP, AI-powered engines, and turnkey frameworks that address most pressing challenges in finance and technology.

Regulatory & compliance acceleration, data enrichment

cut audit and AI costs, and ensure readiness for AI Act, DORA, and beyond. Enrich transactional data.

Auditor™ · Data Enrichment

Integration, lending & payments innovation

an integrated stack of integration services, lending engines, AI underwriting, BNPL flows, and payments accelerators, designed to reduce time-to-market by months and lower operational overhead.

Microloan Origination Engine · SnapCredit AI · UnderwriterAI · Payments Accelerator · Integration Services Accelerator

Next-gen digital banking

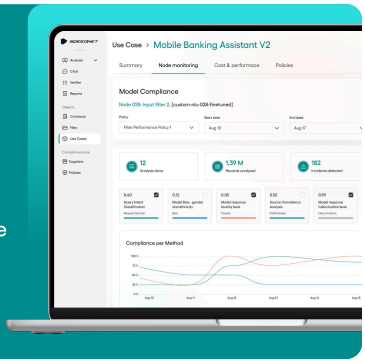
strategic platforms that help banks capture future customers and differentiate their mobile channels with value-added services and superapp-level engagement.

Youth banking platform · Superapp VAS Enablement

SPEEDNET

Auditor™

An AI Management System that lets you manage key AI Governance processes inside the company, monitor and ensure compliance of all AI Systems in real-time.



Delivery acceleration services

Inject speed, flexibility, and engineering excellence directly into your operations.

Our services cluster is designed to break down delivery bottlenecks, modernise architectures, and accelerate digital product development.

Speedflow™

our proprietary AI-assisted engineering framework that accelerates the entire software development lifecycle.

Gemini · GitLab · slack · REDMINE

AI MVP & PoC development

de-risk innovation with rapid prototypes and working demos in weeks, not months.

Agile delivery PODs

cross-functional teams that deliver digital products 4.5x faster while reducing costs.

Compose + KMP LaunchPad

build production-grade mobile apps with 60–80% shared code and native quality.

OpsView 360

unify monitoring and observability to cut incident resolution times by up to 70%.

Velocity pipeline services

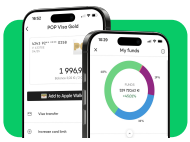
optimise CI/CD pipelines to enable daily releases and accelerate feedback loops.

Working with Speednet means

- Leveraging ready-to-go delivery models and teams, speeding up development without sacrificing quality
- Access to a team that understands BFSI industry and its regulatory compliance
- Flexible scaling - you can increase or reduce capacity as you go
- Technology modernisation, that will keep you up with innovations while staying compliant and secure

Specific experience and knowledge areas

Digital financial products (web & mobile) · AI adoption and AI-driven SDLC · AI-Assisted engineering · APIs and System integration · Mobile banking (retail and coporate) · Payments · Value-Added-Servces · Business logic layer (Middleware) · CMS & CRM systems · Loans and mortgages · Core engines in finance and insurance



The change we delivered for Finnish POP Pankki mobile banking app translated into a **20% increase in customers and a 126% increase in overall profit**, enjoying the highest customer satisfaction rate in the Nordics.

EUVIC

30+ companies

5000+ engineers

We are EUVIC Group's center of banking excellence.

If you're looking for **tangible business results without long implementation cycles**, Speednet is the agile partner for impactful digital transformation at speed.

Ready to start? **Let's talk.**



Adam Rasiewicz
Senior Business Development Manager
adam.rasiewicz@speednetsoftware.com
+48 880 319 710

speednetsoftware.com →