AI RiskWise

AI Governance & Trust Framework

How is AI RiskWise bv tracking full compliance?

V1.0
January 26, 2026
Owner: Head of Compliance

Purpose, Scope & Applicability

This document explains how AI RiskWise designs, builds and operates AI solutions in a responsible, transparent and compliant way. It is written for **procurement, risk, compliance and IT-security teams** evaluating MeetingWise and related services. The scope covers the MeetingWise platform and supporting processes, including data protection, security, AI governance, supplier oversight and customer-facing transparency.

Executive Summary

This **AI Governance & Trust Framework** defines how AI RiskWise BV designs, builds, validates, deploys, and maintains AI systems in a responsible, transparent, and compliant manner. It reflects the company's mission, ethical principles, and commitment to trustworthy AI, while ensuring alignment with applicable regulatory requirements and industry best practices.

The framework is built on six core governance domains:

Model Control
Ensures AI models are developed, tested, validated, monitored, and documented in a controlled and transparent way. This includes lightweight Model Cards, validation steps, and clear disclosure to users when AI is used.

Data Management
Defines how data is collected, processed, stored, and governed throughout the AI lifecycle, including data quality, lineage, access control, and privacy-by-design practices.

Ethics
Embeds our company values — transparency, fairness, and responsible innovation — into product design and decision-making, preventing harmful or unintended outcomes.

Geopolitical & IT Risks
Addresses dependencies on technology providers, jurisdictions, and external infrastructures, including export controls, data-location concerns, and supply-chain resilience.

Legal & Regulatory Requirements
Ensures compliance with applicable regulations such as the **EU AI Act**, the **GDPR**, and contractual obligations. This includes clear privacy notices, terms & conditions, and transparency obligations for limited-risk AI systems.

Cybersecurity
Establishes technical and organizational measures (TOMs) to protect systems, data, and models from unauthorized access, misuse, or compromise, including incident-response procedures and secure development practices.

These six governance domains guide:

- **Product design**
  Every AI component is designed in line with the governance methodology, ensuring transparency, data minimization, and robust security.
- **Website policies**
  Including the Privacy Notice, Terms & Conditions, and the Compliance & Security page, which publicly summarize our approach and obligations.
- **Customer contracts & compliance documentation**
  Such as DPA's, System Cards, security documentation, and other materials customers need to assess us as a vendor.

This framework provides a **cohesive, repeatable, and scalable governance structure** that supports our current low-risk products while preparing the organization for future growth, higher-risk AI applications, and evolving regulatory expectations.

AI Governance & Trust Framework at a glance

The model above illustrates how AI Governance, Business Foundations, Implementation Areas, and Client-Facing Deliverables work together to ensure responsible, transparent, and compliant use of AI within the organization. It shows how high-level strategy, company values, and daily operations are guided by a structured governance methodology and translated into practical product and website implementations. These, in turn, result in a concise set of customer-ready deliverables that provide transparency and build trust.

1. Governance Setup

The AI Governance Methodology forms the backbone of the entire framework.

It is based on AI RiskWise's proprietary governance approach originally developed for AI applications in the regulated financial sector and adapted here for broader use.

The methodology consists of six governance domains:

Model Control – Ensures models are tested, validated, monitored, and documented.

Data Management – Defines how data is collected, stored, processed, protected, and governed throughout the AI lifecycle.

Cyber Security – Establishes technical and organizational safeguards to protect systems, data, and AI outputs.

Legal Requirements – Aligns product design and operations with regulations, such as the GDPR and the EU AI Act requirements for low-risk AI.

AI Ethics – Embeds ethical principles like transparency, fairness, and safety into design and decision-making.

Geopolitical Risks – Addresses supply-chain exposure, vendor dependencies, cloud jurisdictions, and broader IT risk scenarios.

These governance domains provide the structure and guidance for our AI-product development.

2. Business (Strategy, Tactics, Operations)

The business side describes how the organization itself is structured and how governance principles feed into decision-making at all levels.

Strategy: Company Objectives

This level defines the organization's mission, business goals, and structural responsibilities. It also includes the RACI model and assigns ownership of governance-related roles (such as the DPO or security owner).

These objectives set the direction for all AI-related activities.

Tactics: Business Ethics

Business Ethics represent the company's values and expected behaviors.

Core principles, such as transparency, fairness, and safety, act as guiding standards for product development, communication, and risk considerations. They ensure that AI solutions remain aligned with the company's moral and societal expectations.

Operations: Business Operations

This operational layer covers products, websites, and communication activities. It is where governance principles become concrete practices embedded in day-to-day execution, for example in the design of AI features, user messaging, data flows, and documentation.

3. Implementation Areas (Product & Website)

This layer translates governance and business principles into practical execution within the organization.

Product Implementation

Product development must adhere to the AI Governance Methodology. This includes:

- Implementing AI transparency notices
- Documenting the model's intended purpose and limitations
- Ensuring responsible data use and alignment with privacy and security standards

It also includes preparing the materials needed to communicate clearly with clients about how AI works within the product.

Website Implementation

The website serves as a core channel for transparency and compliance. It includes:

- A GDPR-compliant Privacy Notice
- Cookie consent mechanisms
- Clear disclosures about AI usage and limitations

- Terms & Conditions
- A Compliance & Security section summarizing how we govern data and AI

4. Client-Facing Deliverables

The result of governance, business setup, and implementation is a focused set of deliverables that support transparency, trust, and compliance.

For a low-risk AI product, these deliverables are grouped in three compact categories:

1. Contractual & Privacy Essentials

- Terms & Conditions
- Data Processing Agreement
- Privacy Notice
- Subprocessor information

2. Security & Governance Overview

- High-level overview of security practices and safeguards
- Summary of AI Governance methodology (six domains)
- Basic incident-response and data-protection notes
- Overview of how AI is used within the product

3. Product Transparency Package

- System Card (high-level explanation of AI functionality)
- Model Card (purpose, limitations, data profile)
- Datasheet (dataset overview)
- Additional customer transparency notes
- Optional client-specific privacy suggestions

These deliverables provide customers with clear insights into how our systems work, how their data is protected, and how we ensure responsible AI operations.

Summary

This model demonstrates how:

- Governance guides the structure and boundaries of AI development
- Business strategy, ethics, and operations shape how governance is applied
- Product and website implementation put governance into practice
- Client-facing deliverables ensure transparency, trust, and compliance

Together, these components create a coherent, scalable, and auditable framework that enables responsible innovation while meeting regulatory expectations and customer requirements.

Roles and Responsibilities

A clear separation of responsibilities ensures accountability throughout the ecosystem in which MeetingWise is used.

AI RiskWise is responsible for the secure design, hosting and operation of the MeetingWise platform. We manage system security, model quality, data protection, governance documentation, incident handling, availability and updates. We treat customer data solely as processor under the GDPR and only process such data under customer instructions.

Customers remain responsible for informing meeting participants, establishing a lawful basis, reviewing AI-generated outputs, ensuring human oversight by licensed advisers and implementing appropriate internal procedures. Customers also control retention settings and user access.

These responsibilities are mirrored in the Data Processing Agreement and in the Customer Responsibilities Annex to the Terms & Conditions.

Security Governance and Practices

AI RiskWise operates a formal security management framework with defined ownership, documented policies and recurring review cycles. A multidisciplinary team, including our Head of Security, the Head of Compliance and our Data Protection Officer, oversees this programme.

The security programme includes encryption for all data transfers and stored content, strict access control with RBAC, SSO and MFA, detailed audit logging, intrusion detection, secure software development standards, vulnerability assessments and regular penetration testing.

Additionally, we maintain a security roadmap aligned with ISO/IEC 27001 standards. Even before certification, we align our controls with the standard's core requirements so our customers can trust that our processes support audit and assurance expectations.

GDPR Compliance, Data Residency and Retention

We follow a clear and simple data protection model:

- The customer is the controller,
- AI RiskWise is the processor,
- All processing takes place entirely within the EU,
- Audio is deleted immediately after processing,
- Transcripts and summaries are retained for up to 60 days (30 by default),
- Technical and security logs are retained for up to 365 days.

We never train models on customer data. We provide full transparency on processing purpose, legal basis, rights and security commitments. Data subjects can exercise rights through their controller; we support controllers through our DPA and established internal procedures.

AI Model Lifecycle, Monitoring and Escalation

Our model lifecycle includes stages for design, evaluation, deployment, monitoring and re-assessment. Prior to deployment, models undergo testing to check functional correctness and to identify potential risks or limitations. After deployment, we monitor for drift, errors and quality deviations.

If model behaviour deviates from expected results or poses potential harm (for example, misleading or incomplete summaries), we treat this as an AI incident. Incidents follow a defined flow: detection, triage, containment, customer communication where appropriate, remediation and documentation. This approach ensures transparency, responsiveness and continuous improvement.

Geopolitical and IT Risk Mitigation

Our EU-only hosting strategy reduces exposure to foreign jurisdictions and to disruptive regulatory mismatch. All compute, storage and AI inference are carried out exclusively in EU data centres provided by our service providers / subprocessors.

We maintain contingency plans for supplier disruption, regularly assess the legal and operational stability of our vendors, and ensure that contracts include EU-adequacy, confidentiality and security commitments.

Vendor and Subprocessor Governance

We carefully select and monitor our subprocessors. Before onboarding a provider, we perform due diligence that includes security evaluations, privacy compatibility checks, data-location assurances, operational resilience assessments and contractual review.

During the relationship, we evaluate providers on a recurring basis and monitor their regulatory posture, uptime, performance and security posture. Customers are notified in advance of any intended subprocessor changes, and concerns are addressed cooperatively.

Product and Website Implementation

Our governance model is reflected in both product design and our website disclosures.

In the product, governance principles appear as privacy-by-design choices, clear AI explanations within workflows, reliable data handling practices, scoped permissions and transparent retention controls. Documentation describing model purpose and limitations is available to customers as part of our transparency commitments.

On the website, we provide GDPR-aligned privacy notices, cookie consent mechanisms, Terms & Conditions, information about AI use and limitations, and a Compliance & Security section summarising our governance approach. These materials support customers' need for clarity, audit readiness and compliance alignment.

Deliverables Provided to Customers

To support due-diligence processes, AI RiskWise provides a standardised set of documents designed to be concise yet comprehensive. They include:

- Contractual and privacy essentials such as the Terms & Conditions, DPA, Privacy Notice and subprocessor information.
- Security and governance summaries offering an accessible overview of our technical safeguards, governance principles, incident response approach and system behaviour.
- A Product Transparency Package containing a System Card, Model Card (light), dataset information, and explanations of AI limitations and controls.

Business Continuity and Disaster Recovery

We maintain continuity across our operations through documented BCP and DRP measures.

Our strategies include secure data backups, restoration testing and defined recovery targets (RTO/RPO available under NDA). These measures are tested periodically to ensure robustness and fast recovery in case of an outage or unexpected disruption.

This guarantees that customers can rely on MeetingWise for the continuity of their advisory workflows.

Security Incident Response

Our incident response procedures ensure rapid and coordinated action when a security event occurs. Incidents are monitored through 24/7 automated systems combined with operational oversight.

When a personal data breach is detected, we notify affected customers without undue delay and aim to provide relevant details within 72 hours, including the nature of the incident, its impact and appropriate containment steps.

We cooperate with customers to support any regulatory notifications they may need to perform and we conduct internal reviews to prevent recurrence.

Regulatory Alignment and Sector Expectations

MeetingWise is designed to support organisations operating under the IDD and Wft frameworks, as well as best-practice expectations from supervisory bodies. The

system enhances documentation quality, supports advisers in maintaining audit-ready client files, and increases operational consistency — without ever taking over decision-making.

Under the EU AI Act, MeetingWise is classified as a limited-risk system. We meet the transparency requirements by clearly explaining when and how AI is used.

Our practices also follow the principles of the DORA regulation by applying structured resilience measures, incident management and vendor governance.

The AI RiskWise Trust Pack

To simplify procurement and compliance reviews, we provide a curated “Trust Pack” containing:

- The Terms & Conditions
- Data Processing Agreement
- Privacy Notice
- SLA-Light
- This Governance & Compliance Overview
- Subprocessor overview
- The AI Product Information Document (AIPID)
- MeetingWise System Card
- Model Card (Light)
- Dataset Datasheet (Light)

This helps organisations complete vendor assessments quickly and with confidence.

Change Management and Document Versioning

AI RiskWise maintains a change log for governance materials and will proactively notify customers of material changes that affect commitments or core functionality. Historical versions and summaries of updates are available upon request.

Contact

For questions about governance, security, privacy or compliance, please contact:

AI RiskWise BV

A. de Bruynkade 38

3059 SE Rotterdam

The Netherlands

Email: legal@airiskwise.com

Website: www.airiskwise.com