

Mastek's Approach to Secure and Responsible AI

SECURE • PRIVACY-CONSCIOUS • ETHICAL • ACCOUNTABLE

Contents

- 1** Preamble
- 2** Scope
- 3** Responsible AI Policy Objective
- 4** Commitment from top management
- 5** Responsible AI Principles
 - 5.1** Privacy and responsible use of data
 - 5.2** Securing the AI Ecosystem
 - 5.3** Fairness, transparency and human oversight
 - 5.4** Our people and our partners
 - 5.5** Training
- 6** Updates to this policy

1. Preamble

Artificial intelligence is reshaping how organisations serve their customers and run their operations. At Mastek, we are committed to embracing this opportunity responsibly adopting and developing AI in a way that is designed to be secure, privacy-conscious, ethical and accountable.

2. Scope

It is intended to apply to our people and, where relevant, to the consultants, contractors and third parties who work with us, as well as to the AI-enabled tools, platforms and services we own or operate.

3. Responsible AI Policy Objective

This statement provides a high-level overview of the principles that guide the responsible use and development of AI across Mastek. It sets out our approach to ensuring that AI is designed, deployed and managed in a secure, privacy-conscious, ethical and accountable manner, while supporting innovation and compliance with applicable legal and regulatory requirements.

4. Commitment from executive management

Our AI activities are overseen through an internal governance forum that brings together security, privacy, technology and business leadership, and that reviews higher-risk AI initiatives and considers any exceptions to our internal requirements. Mastek's Risk Management and Governance Committee provides oversight of enterprise risks, including those related to Artificial Intelligence.

We maintain a governance framework that seeks to address AI across its lifecycle, with responsibilities defined for our people, our security leadership and our internal AI centre of excellence. AI systems are assessed for information security and data-related risk and boundaries are set before they are adopted, developed or deployed, and we aim to understand the data our AI systems use, including how models and algorithms we develop are stored and reused. AI systems in operation are subject to ongoing risk oversight, and findings and associated actions are recorded and tracked through an internal AI risk register.

We maintain an internal list of approved AI platforms, and the use of AI tools outside that list is not permitted without appropriate review and approval. Changes to AI technology follow our established change management processes, and we address non-compliance with our internal requirements through the appropriate internal channels.

5. Responsible AI Principles

5.1 Privacy and responsible use of data

We seek to apply the principle of data minimisation, using and sharing only the data reasonably necessary for a defined AI use case, with the purpose of data collection, processing and reuse documented for the AI systems we operate. Where the law requires consent for the use of personal or otherwise protected data, we seek to obtain it.

AI-related data is protected using industry-recognised encryption in transit and at rest, and we define retention periods for AI training and operational data together with secure deletion practices, which are intended to extend to models that may have incorporated personal data. Privacy impact assessments are carried out for AI systems that process personal data and are reviewed periodically and on significant change.

Our data classification framework governs what information may be used with AI tools. Limiting access to sensitive AI capabilities include highly sensitive information, such as personal data, health information, payment data and credentials is restricted from use on external AI platforms, confidential information requires prior approval together with masking or anonymisation where appropriate, and identifying details are expected to be removed or replaced before information is submitted to AI tools.

5.2 Securing the AI Ecosystem

AI systems are held to the same security discipline as the rest of our technology estate. Our practices are designed to include secure coding, peer review and quality gates during development, together with version control over system code and training data. We aim to use trusted and appropriately validated datasets within secured environments, and we apply measures intended to identify and reduce bias and inaccurate or unreliable model output.

Deployment is managed through controlled pipelines with access controls and audit logging, production environments are appropriately segregated, and supporting infrastructure is maintained and patched in line with our standards. Access is governed by role-based access control and least-privilege principles, supported by multi-factor authentication and time-bound access for administrative functions. Application programming interfaces are authenticated, authorised and managed through controlled gateways, and approved platforms are accessed through enterprise-managed accounts with logging and monitoring enabled. AI-related security events are managed through our established incident management process, including notification to regulators, authorities and affected parties where applicable law or contract requires.

5.3 Fairness, transparency, accountability and human oversight

We apply processes intended to identify and mitigate potential bias across the AI lifecycle, and, where relevant to the use case, we assess AI systems for fairness. Where explanation is appropriate or required, we work to provide meaningful information about how AI-assisted decisions are reached, and we maintain documentation of known model limitations and associated risks.

Human in the Loop approach: Human oversight remains central to our approach. Our people are expected to be accountable and review AI-generated output for accuracy, fairness and appropriateness before it is relied upon in business decisions, communications or operational processes. We set clear boundaries on acceptable use: AI must support legitimate business objectives, respect applicable law, and protect privacy and security. AI must not be used to obtain unauthorised access to data or systems, to unfairly discriminate against individuals or groups, to conduct malicious or offensive cyber activity, or to process personal data without an appropriate lawful basis. We are committed to conducting our AI activities in a manner consistent with applicable legal and regulatory requirements, including data protection and intellectual property obligations, and we continue to track developments in AI-specific regulation.

5.4 Our people and our partners

We provide periodic awareness and education for our people covering AI ethics, fairness and non-discrimination, data privacy, data quality and curation, intellectual property and security. Third-party AI providers are selected on their ability to meet our business and assurance expectations and are subject to security, privacy and data-use assessments. Security and privacy requirements are reflected in our contractual arrangements where appropriate, audit rights are sought for critical AI services, and third-party AI services are monitored and periodically reviewed.

5.5 Training

We incorporate a comprehensive security awareness program combining mandatory, tool-based enterprise-wide training across Information Security Management System (ISMS) and Responsible AI Usage. These trainings are instituted at the time of employee and contractors onboarding and reinforced through mandatory annual refresher programs to ensure sustained awareness and compliance. Training completion is centrally tracked and monitored to ensure full participation, with compliance status regularly reviewed and reported.

6. Updates to this policy

Responsible AI is an evolving discipline. We keep our approach under review and continue to strengthen it in response to emerging threats, regulatory developments and advances in technology. Any deviation related to this policy must be formally approved by Cyber Security, Privacy and AI Forum.