TECH mahindra

Responsible Artificial Intelligence (AI) Program

Contents

Introduction 3
Governance And Accountability 3
Oversight, Risk Management, and Model Assurance 3
Continuous Learning and Capability Building 5
Legal and Regulatory Compliance 5
Ethical Deployment and Oversight 6
Ethical AI 6
Transparency and Labeling of AI-Generated Content 6
Human-Centered AI Design 7
Customer Transparency 7
Access Controls for Sensitive AI Capabilities 7
Environmental Sustainability of AI 8
AI-Driven Innovation and Platform 9
About Tech Mahindra 1

Introduction

Artificial Intelligence (AI) is no longer a project inside the enterprise. It is becoming an enterprise. The shift from AI experimentation to scaled AI execution has changed the question organizations need to answer. Not whether to adopt AI, but whether AI can be adopted at speed without compromising on trust, fairness, security, or sustainability. These are not soft considerations. They are the conditions under which AI creates durable value rather than short-term advantage that erodes. Regulatory frameworks are tightening around them, customers are selecting partners on the basis of them, and the individuals whose work and lives are touched by AI systems are asking harder questions about how those systems were built, what data trained them, and who is accountable when they get something wrong.

For Tech Mahindra, these questions are an operating reality. AI Delivered Right is the name we give to the discipline of answering them — four pillars, executed in parallel, built on the conviction that responsible governance and business performance are not trade-offs. They have the same objective.

Transformation Delivered Embedding AI into core business processes to unlock new revenue models, customer experiences, and operating constructs — not just point automations.	**Productivity Delivered** Driving measurable efficiency through agentic automation, intelligent decisioning, and AI-led delivery across software engineering, operations, and BPS.
Innovation Delivered Building new products, services, and IP — including domain-specific language models and industry agents — that move clients ahead of their markets.	**Assurance Delivered** Embedding responsible AI, security, governance, and explainability into every deployment, from design through production.

Governance And Accountability

Oversight, Risk Management, and Model Assurance

Effective AI governance at Tech Mahindra is anchored in strong oversight, accountability, and risk management across the entire AI lifecycle to ensure systems operate responsibly, securely, and in alignment with legal and ethical expectations. Recognizing that AI systems operating in dynamic environments can experience shifts in performance over time, known as model **drift**, and may occasionally generate outputs that are plausible yet inaccurate, referred to as **hallucinations**, we have established robust controls to actively manage these risks. Through continuous monitoring, structured review mechanisms, and human oversight, we ensure AI outcomes remain reliable and aligned to

intended objectives. Clear governance structures and grievance redressal mechanisms further strengthen transparency, explainability, and stakeholder trust.

- We have embedded governance, security, and oversight mechanisms across the AI lifecycle to support the responsible deployment and controlled use of AI capabilities. Through our Assurance Delivered approach, responsible AI, security, governance, and explainability are integrated from design through production, helping ensure that higher-risk and sensitive AI applications are subject to appropriate controls, monitoring, and review in accordance with legal, ethical, and regulatory requirements.

- Associates are required to maintain appropriate human oversight throughout the AI lifecycle and ensure that AI-generated outputs are reviewed and validated before use. Human-in-the-loop practices, supported by continuous monitoring and periodic validation of AI outputs against expected parameters, help reduce the risk of harmful responses, AI hallucinations, and model drift. Where required, identified risks or deviations are addressed through established governance processes, including additional validation, controlled model retraining or parameter adjustment, version rollback where appropriate, and realignment with defined guardrails and approved usage requirements, along with addressing intellectual property concerns and data privacy risks, while supporting the reliability, accountability, and responsible use of AI-enabled solutions.

- Our AI systems are developed, operated, and governed within a management framework that is externally verified and subject to rigorous, independent third-party assessments against multiple internationally recognized standards. Software and service development processes, including those underpinning AI model development, integration, and delivery, are externally appraised at CMMI Development and Services Version 3.0 Level 5, the highest achievable maturity level under the Capability Maturity Model Integration framework, providing independent validation that TechM's development and delivery practices are disciplined, measured, and continuously improved. The security governance of AI systems is independently certified under ISO 27001:2022, covering controls over data handling, access management, logging, incident response, and the secure operation of AI-enabled environments. Privacy governance applicable to AI systems processing personal data is independently certified under ISO 27701:2019, confirming the implementation of privacy-by-design principles across AI-relevant data management processes. Together, these externally assessed and certified management systems provide third-party assurance across the development, security, and privacy dimensions of TechM's AI management environment.

- We have established clear, accessible channels through which any associate, partner, customer, supplier, or community member may raise concerns, seek clarification, or formally contest outcomes related to AI-driven solutions and decisions. Where an individual or organization is affected by an AI-enabled outcome, they may use these channels to request review, seek explanation of the factors informing that outcome, or escalate concerns for further investigation. All reports are handled anonymously where preferred, through independently managed, multilingual platforms under strict confidentiality protect.

Online Reporting CorporateOmbudsman@techmahindra.com	**Toll-free 24×7 Helpline** India: +91 120-488-4450 US: 000911204884450
Offline Reporting Sudeep Chopra, Corporate Ombudsman, Tech Mahindra Ltd, Plot no 58A and B, NSEZ, Phase II, Noida (UP) 201301, India	**Quick Online Reporting (Associates)** Available via the SOS button on the Twingo intranet homepage.

Continuous Learning and Capability Building

At TechM, continuous learning underpins the responsible development and use of AI. We invest in ongoing research, structured training, and capability-building initiatives to strengthen AI-related competencies across the workforce while addressing evolving risks and expectations.
Our approach integrates technical skill development with awareness of emerging challenges, including biases, errors, and security and privacy risks. Through continuous monitoring and improvement, we enhance the reliability and accountability of AI-enabled capabilities and support their responsible adoption across business contexts.

Our training programmes are designed to equip associates with a comprehensive understanding of AI use in practice, covering its ethical implications, bias awareness, data protection requirements, and security considerations. This includes structured modules on recognizing and mitigating algorithmic bias, responsible data use, intellectual property considerations in AI-generated work, and the security responsibilities that accompany operating AI systems at enterprise scale. Training is delivered through a range of structured pathways including AI belt certification tracks, GenAI and GitHub Copilot certification programmes, and mandatory cybersecurity and data privacy awareness modules aligned with our broader Cybersecurity and Data Privacy commitments. This enables associates to make informed decisions, apply appropriate judgement, and ensure AI is used in a manner that is responsible, secure, and aligned with organizational and ethical standards.

We continue to drive workforce transformation through targeted reskilling and upskilling initiatives, enabling our teams to adapt to evolving AI-driven roles and deliver value across client engagements. This focus on capability building has strengthened enterprise-wide readiness and accelerated the integration of AI into solution design and delivery.

45,000+	~80%	8	84%
Associates trained/certified in AI (FY26)	Associates AI-enabled through learning	AI & Advanced Analytics Labs	Client-facing associates on AI-enabled projects

Legal and Regulatory Compliance

The regulatory landscape governing AI is developing rapidly and unevenly across the geographies in which we operate. We actively monitor the evolution of AI-related laws and regulations, including the EU AI Act, relevant US state-level requirements, and emerging frameworks across Asia-Pacific, to

ensure our AI-enabled solutions are designed and deployed within applicable legal boundaries. This includes ongoing alignment with data protection, privacy, and security requirements, as well as industry-specific regulatory obligations relevant to our clients' sectors. Our internal governance practices, model documentation standards, and risk-assessment processes are continually updated to reflect this evolving landscape, supporting our clients' compliance needs as well as our own.

Ethical Deployment and Oversight

As AI becomes increasingly integrated into business operations, organizations must effectively manage risks related to bias, fairness, privacy, security, transparency, and regulatory compliance. AI systems may experience challenges such as model drift, where performance declines due to changing data patterns, or AI hallucinations, where inaccurate or misleading outputs are generated. Continuous monitoring, validation, and human oversight are therefore essential to maintain the reliability, safety, and effectiveness of AI-enabled solutions.

Ethical AI

We have implemented ethical AI principles into the development and deployment of AI-enabled solutions by promoting fairness, impartiality, and accountability in AI-driven outcomes. Deployed AI models are subject to structured assessments covering accuracy, consistency, and fairness, conducted through TechM VerifAI's analytics and observability layer and reviewed by designated governance stakeholders. These assessments evaluate model outputs for potential bias, demographic disparities, and unintended impacts against expected parameters, with findings reviewed and remediation actions assigned where deviations are identified. Human oversight, risk assessment, and validation mechanisms further support the reliability of AI applications and ensure they remain aligned with approved use cases, data privacy requirements, and organizational standards.

Inside our Makers Lab, Responsible AI operates as a dedicated research and development pillar. The mandate is to advance frameworks for fairness, transparency, and accountability in AI systems, ensuring deployed models are both explainable — capable of surfacing the reasoning behind a decision — and resilient to bias. Bias detection and mitigation pipelines are embedded within our Responsible AI governance protocols, with the Makers Lab Responsible AI pillar advancing these frameworks across GenAI and Continual AI systems to ensure models are evaluated for fairness and bias resilience as part of the governance review process, and explainable AI dashboards are being developed to give stakeholders direct visibility into how a decision was reached. Responsible AI principles from this pillar are applied in education, healthcare, and enterprise automation use cases with explicit focus on inclusivity, multilingual accessibility, and equitable outcomes.

Transparency and Labeling of AI-Generated Content

We have embedded transparency within our AI-enabled solutions by providing appropriate information on how AI technologies operate, including their limitations, potential biases, and intended use where relevant. The Company ensures clear identification and labelling of AI-generated content, including text, images, and videos, as well as outputs or decisions supported by AI systems, helping users distinguish AI-generated or AI-assisted information from human-created content. Users are proactively

informed when they are interacting with an AI system rather than a human, enabling informed choices and strengthening trust, accountability, and transparency across AI-assisted interactions.

Human-Centered AI Design

AI at Tech Mahindra is designed to augment human capability, not replace human judgement. Core professional disciplines, including software development, remain directed by our people, with AI deployed to support and accelerate this work rather than assume it. We actively seek and incorporate user feedback throughout the design and development of AI-enabled solutions to ensure they remain relevant, usable, and aligned to real operational needs, while encouraging associates to apply AI in ways that reduce routine effort and create space for more meaningful work.

Customer Transparency

We are transparent with our customers about the nature of AI-enabled services, including their capabilities and limitations, so that expectations are set clearly from the outset. Where customers direct the use of specific AI tools or technologies in their engagements, the terms of that arrangement, including the allocation of associated risk and responsibility, are agreed and documented as part of the commercial relationship. This approach ensures clarity for our customers and supports responsible, well-governed use of AI across client engagements.

Access Controls for Sensitive AI Capabilities

We have implemented strict access control to ensure that the use of sensitive AI capabilities occurs within approved purposes, defined operational boundaries, and established oversight mechanisms. These measures include restricting access to authorized users, defining responsibilities, and applying safeguards, control points, and escalation mechanisms to help prevent unauthorized use, mitigate ethical and privacy risks, and support compliance with applicable legal, regulatory, and organizational requirements. Controls are specifically designed to prevent application to unauthorized biometric surveillance, social scoring, misuse, unintended application, or any use case outside the approved scope. Prior to deployment of any sensitive AI capability, a formal review process is conducted to assess the intended use case, applicable data protection requirements, legal compliance considerations, and oversight mechanisms, with approval required from designated governance stakeholders before the capability may be activated in any production environment.

Access to AI applications and tools is further governed through approved enterprise licenses, secure workspaces, Tech identities, authentication requirements, and prior review of AI tools, platforms, components, and vendors before use. This helps ensure that sensitive AI capabilities are deployed only within permitted use cases, documented limitations, and defined data protection and security boundaries.

For AI solutions in regulated, high-sensitivity environments, these boundaries are structurally enforced rather than only stated. A flagship public-sector engagement in FY26 required full data sovereignty: off-the-shelf, public-cloud AI was ruled out entirely. The solution — combining orchestration, persistent memory, vector-based document intelligence, and live system integrations — was delivered to operate entirely within a sovereign infrastructure footprint, with no exceptions on data residency, access control, or audit trail. Vendor selection for AI technology is governed by the same discipline: AI tool

vendors are evaluated on their disclosures of model construction, training data provenance, reliability, and intellectual property risk, not only on functional capability.

Environmental Sustainability of AI

As AI adoption grows, we recognize the need to manage not only its social and economic implications but also its environmental footprint. AI systems can require significant compute capacity, energy, and infrastructure support, which may contribute to environmental impacts if not managed responsibly. Therefore, guided by Tech Mahindra's AI philosophy, we aim to minimize the environmental impact of AI through energy-efficient design, responsible infrastructure choices, and sustainability considerations across the technology lifecycle. This aligns with our broader commitment to resource efficiency, carbon reduction, and responsible innovation.

We have integrated green data centre principles through energy-efficient infrastructure, renewable energy integration, advanced cooling technologies, and continuous monitoring of Power Usage Effectiveness. These measures help optimise energy consumption, reduce carbon emissions, improve resource efficiency, and support our broader sustainability commitments. Average PUE has improved from 1.96 in FY23 to 1.73 in FY26; renewable energy share at TechM data centres has grown from 1.77% at baseline to 35.66% in FY26.

Our Green IT approach operates on two complementary fronts. **Green WITH Software** uses cloud computing, AI, and data centre technology to reduce enterprise-wide carbon emissions and improve resource utilization. **Green WITHIN Software** embeds eco-friendly coding practices, energy-efficient algorithms, and responsible AI design directly into the development lifecycle so sustainability is built in from the first line of code, not measured after the fact.

Four programmes operationalize this approach across the organization:

Green CodeRefiner™ (GCR™) AI-powered code carbon-optimisation utility that analyses application source code, identifies inefficient patterns through dynamic carbon emission analysis, and automatically transforms energy-intensive code segments into more efficient alternatives. Works across multiple programming languages and integrates directly into CI/CD pipelines. Measured result: 30–40% improvement in Green Impact Scores.	**Green Cloud Carbon Footprint (GCCF)** Monitoring utility that tracks cloud emissions across AWS, Azure, and GCP, offering detailed visualizations by region, service, and resource type to support greener cloud operations. Enables delivery teams to act on actual consumption hotspots rather than managing against aggregate estimates.
Green IT AI-Powered Assessment A structured framework that evaluates IT sustainability maturity, identifies emission hotspots across infrastructure and	**Green IT Operations Dashboards** Custom dashboards provide continuous visibility into emissions across workloads and infrastructure, enabling teams to identify

applications, and provides actionable roadmaps for emission reduction and energy optimization. Supports clients in building baseline visibility before targeting reductions.	hotspots, track reduction progress, and incorporate sustainability into operational decision-making as a routine practice.

We track and assess the environmental performance of AI and Green IT initiatives using defined metrics and KPIs, including energy efficiency, carbon reduction, cloud emissions visibility, and Green Impact Scores. This data-driven approach enables quantification of sustainability outcomes and supports continuous improvement in how AI-enabled tools contribute to lower energy consumption, improved resource efficiency, and reduced emissions. Green CodeRefiner™ provides a clear AI-attributed sustainability metric through its measured **30–40% improvement in Green Impact Scores**, demonstrating how AI can be used not only as a business accelerator but also as an enabler of sustainable technology operations.

AI-Driven Innovation and Platform

We are actively building new AI products, platforms, and capabilities designed for scale, responsibility, and measurable outcomes. Project Indus, our family of proprietary foundation and domain-tuned models, includes a co-innovation initiative with Qualcomm and Intel to develop a Hindi-first education LLM evolving into a civilization-scale Indic-language foundation model supporting Hindi and key dialects. This initiative supports sovereign AI capabilities, localized AI deployment, digital inclusion, and efficient edge-based AI solutions that create long-term value for communities, businesses, and society.
We have also deployed TAN — The Ability Network, a co-innovation platform developed in partnership with Tech Mahindra Foundation, which facilitates anonymous connections between individuals with disabilities and appropriate caregivers. TAN demonstrates how responsible technology deployment can advance equitable access, dignity, and positive social outcomes at scale.
We have implemented TechM Orion, an enterprise-grade agentic AI platform, to support AI-led transformation through the responsible orchestration of complex business workflows. By embedding governance, risk management, and scalability considerations into AI-enabled operations, the platform helps deliver secure, accountable, and trustworthy outcomes across cloud, on-premises, and hybrid environments. Orion's built-in governance layer, which maintains audit trails of AI-generated outputs and agent interactions, providing a verifiable log that labelling and disclosure requirements have been applied consistently across AI-assisted services and workflows.
TechM VerifAI, part of our Orion and Indus platform stack, provides dedicated safety, compliance, and performance assurance through monitoring tools for observability, versioning, and analytics in deployed AI systems. VerifAI enables continuous visibility into model performance after go-live, supports detection of performance degradation and compliance drift, and facilitates controlled retraining and versioning.

About Tech Mahindra

Tech Mahindra (NSE: TECHM, BSE: 532755) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 146,000+ professionals across 90 countries, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognizes global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies.