Policy
Artificial Intelligence & Machine Learning

Introduction

Our vision is that we build human centered Ai-powered solutions, which enable faster smarter decisions to create value for our communities.

We will deliver:

- Community Value – our deep understanding of community needs drives effective solutions
- Actionable Insights – data analytics and insights to drive smarter, faster decisions
- Ethical AI – a secure unified AI platform used for all solutions

This will be achieved through:

Incremental Adoption

We are integrating AI and machine learning into key areas of our product suite, balancing innovation with necessary safeguards to address risks.

Risk-based Approach

We recognise varying comfort levels with AI applications. We are implementing guardrails to reduce the likelihood of harm occurring in high-risk settings.

Continuous Evaluation

We actively research and invest in new AI tools to effectively and elegantly address customer business challenges, prioritising maturity, reliability and security.

Data & Governance

We constantly invest in evolving our Ai and data governance framework that guide our internal teams to design, build and deliver reliable, transparent and safe AI scenarios

Purpose

The purpose of this document is to provide guidelines and principles for the responsible, safe, ethical and secure development of AI/ML powered ERP solutions and to be transparent in our governance approach. Importantly, we continuously evaluate this AI/ML approach and our governance framework, prioritising maturity, reliability and security

Scope

This policy applies to all employees, contractors and third parties involved in the development, deployment and support of TechnologyOne's AI/ML powered solutions.

Guiding AI Principles

1. Create Value

TechnologyOne's AI/ML powered solutions will create value for our teams, customers, and the communities we serve. Our AI/ML-powered ERP solutions will:

- Make a tangible difference to our stakeholder's business practices by enhancing processes and supporting efficiency
- Focus on repetitive & low-value tasks, thus enabling stakeholders to perform higher-value work.
- Be resilient, sustainable, and able to withstand unexpected adverse events.
- Be tested for validity and perform reliably over the life of the system, and in line with our commitment to safety, security & quality

2. Create Collaboratively

We will take a considered, incremental approach to adopting AI/ML by consulting and collaborating with our stakeholders and customers, through proof of concept (POC) projects and early adopter (EA) programs.

Together we will ensure that AI-powered AI solutions:

- Are inclusive, fair, and accessible, ensuring the risk of harmful bias is mitigated at all stages including through development, optimisation, and operational stages
- Incorporate accountability and transparency, including indicating when AI powered functionality is in use.
- Provide traceability of data, processes, and decisions across all stages of tool use
- Provide the ability for customers to choose to use AI powered functionality, in line with their own organisational needs, values and context.
- Empower end users to self-serve, where possible, to manage the consumption and configuration, and addressing exceptions or errors.

3. Create Empathetically

Maintaining trust and meeting the expectations of our stakeholders is fundamental to successfully bringing our AI vision to life. Creating AI-powered solutions empathetically means we provide a range of technologies to enhance the way problems are solved, but in a respectful, considered way.

Our AI-powered solutions will:

- Avoid full automation in areas where human input is required or recommended to set rules, interpret results, or make critical decisions.
- Be created with integrity, ensuring tools are secure and safe, and comply with current and emerging AI, security, and privacy related laws.
- Enable our customers to comply with laws, standards and codes relevant to their business.
- Be risk-assessed prior to deployment and throughout the tool's lifecycle.
- Be designed to be robust, sustainable and perform as expected.

Transparency

The following sections provide transparency into our governance processes and use of AI:

Usage Patterns

Our software has incorporated AI/ML in the following ways:

- Workplace productivity – use of AI/ML tools to streamline ERP processes such as automated information extraction, summarisation, recommendations, natural language searches and a virtual assistant.
- Analytics for insights – use of AI/ML tools for data visualisation via natural language processing, forecasting and prediction.
- Image processing – processes images to automatically identify patterns and objects such as faces, buildings and objects.

Domains of AI Usage

TechnologyOne's application of AI/ML within our software includes:

- Service delivery - enhances efficiency or accuracy of business processes by providing tailored and responsive services. This may include in direct interaction with the public, such as chat-bots, enhanced customer self-service and multilingual capabilities, or support staff or systems which deliver services
- Corporate and enabling - supports corporate functions including HR and finance by automating processes, optimising resource allocation and improving operational efficiency.
- Compliance and fraud detection - identifies patterns or anomalies in data to detect fraudulent activities and ensure compliance with laws and regulations.

AI/ML Governance

TechnologyOne's AI governance framework is aligned to ISO/IEC 42001, which fully supports and operationalises the principles of governance, risk management, transparency, testing and human oversight. Our governance framework is designed to encompass guidance from specific jurisdictions such as Australia, New Zealand and the United Kingdom.

Our AI management system includes but is not limited:

- **AI governance framework and accountability**– TechnologyOne maintains a structured governance framework to oversee AI initiatives across their lifecycle including development, deployment and data management. Executive oversight is provided by our Chief Technology Officer (CTO) and supported by our R&D Leaders and internal governance committee. This ensures AI/ML initiatives are managed consistently and in line with defined principles, policies and risk tolerances.
- **AI risk management** – identification, assessment, treatment and monitoring of risks associated with AI/ML.
- **AI system impact assessments** – evaluating the potential impacts of AI/ML systems on individuals, groups and society to inform design and risk decisions
- **AI system design and development** – designing and developing AI-enabled software in accordance with our AI principles, safety, security and quality standards

- **Testing, monitoring & improvement** – testing our AI enabled-software prior to deployment and monitoring performance over time to ensure it performs as intended and expected outcomes are delivered
- **Stakeholder engagement, transparency and redress** – proactive engagement with customers and stakeholders through clear communication and appropriate mechanisms to support transparency, accountability and responsible AI use.

AI/ML Use Case Classifications

To ensure TechnologyOne's use of AI/ML tools aligns with the principles above and TechnologyOne's appetite for risk, internal teams assess proposed AI/ML use cases against a set of criteria, as summarised below:

Green Path: Proceed with caution

- Use case aligns with the guiding principles
- Use case maintains our security & compliance posture
- Use case aligns with strategic goals
- Use case is feasible

If the use-case meets the criteria above, it can proceed to the next stage of development, in line with our internal development methodology.

Amber Path: Review, adjustment and/or consultation required

- Alignment with guiding principles is unclear or not all can be met
- Feasibility is uncertain
- Considered a 'high risk' model under AI/ML laws, directives, regulations, or standards

'Amber' use cases must address misalignments and modify the design to return it to the 'Green' path. Where this is not possible, the CTO and/or RCS Council may approve an exception, with subject matter expertise input, to proceed.

Red Path: Do not proceed

- Use case is not aligned with the guiding principles
- Use case is not feasible
- Considered a 'prohibited' model under AI/ML laws, directives, regulations, or standards

Where a proposed use case is on the 'Red' path, it cannot proceed unless there are material adjustments made in the use-case design or approach.

AI/ML Shared Responsibility Model

To ensure TechnologyOne's AI powered ERP solutions are trusted, secure, and safe to use, all parties will have a role to play in protecting the security and privacy of the data and systems used. The below diagram provides a summary of the various stakeholders and their responsibilities, in the context of TechnologyOne's AI powered solutions. For more information regarding TechnologyOne's general responsibilities as a SaaS provider, customers are encouraged to request our SOC 2 report

AI/ML Shared Responsibility Model

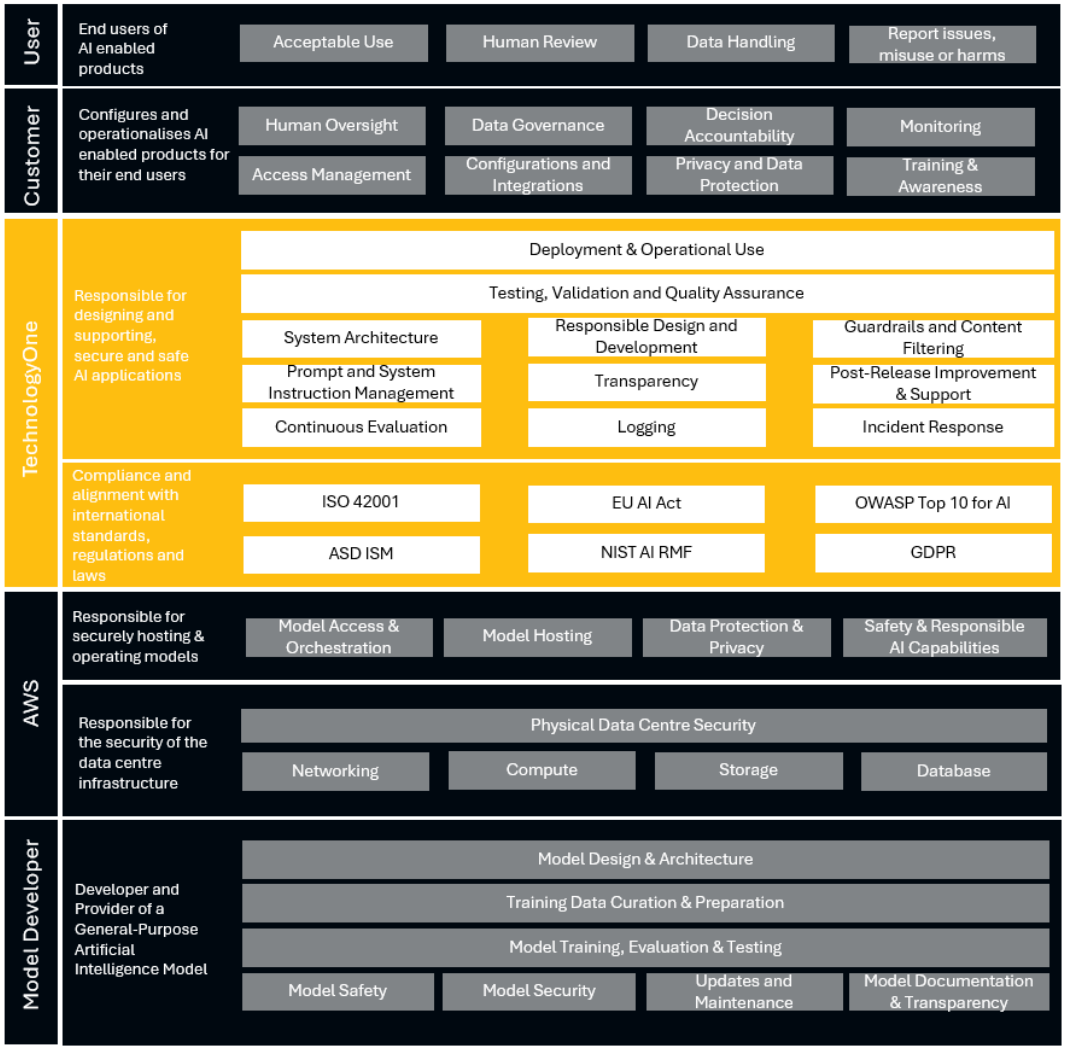

Further details of each set of responsibilities are provided below:

User

Users are responsible for the safe and appropriate use of AI-enabled products to support reliable outcomes. Specifically:

- Using AI products only for permitted, lawful and intended purposes
- Reviewing AI outputs before use or actions
- Inputting only permitted and authorised data

Customers

Customers are responsible for configuring and operating AI-enabled products within their organisation using appropriate controls, oversight, and governance to support safe and effective use. This includes:

- Managing user access to AI-enabled products
- Configuring AI-enabled products in line with business requirements and product guidance
- Controlling data model access in accordance with internal data handling and classification requirements
- Providing appropriate training and awareness aligned to organisational policies and product guidance
- Applying ongoing human oversight to AI-assisted activities and workflows
- Governing data use, privacy, and protection in accordance with applicable requirements
- Maintaining accountability for decisions informed by AI outputs
- Performing appropriate due diligence and testing
- Maintaining policies and procedures to govern the use of AI, data, applications, and systems, and monitoring compliance
- Reporting incidents and issues in a timely manner via the Customer Community
- Ensuring compliance with applicable legal, regulatory and contractual obligations

TechnologyOne

TechnologyOne is responsible for designing, developing, deploying and supporting AI enabled products in a manner that promotes security, safety, reliability and alignment to frameworks and standards. This includes:

- Designing and developing AI capabilities with responsible design principles and built-in safeguards
- Establishing and maintaining secure system architecture to support AI workloads
- Managing prompt and system instruction controls where applicable
- Performing testing, validation, and quality assurance throughout the AI lifecycle
- Managing components used in the AI supply chain
- Implementing guardrails, content filtering, and other safety mechanisms
- Providing transparency through appropriate documentation and design practices
- Enabling logging and monitoring to support oversight and investigation
- Conducting continuous evaluation and post-release improvement
- Providing operational support and responding to incidents related to AI services
- Ensuring ongoing compliance with contractual terms, regulatory requirements, and applicable AI, privacy, and data protection obligations

Amazon Web Services (AWS)

TechnologyOne leverages AWS to deliver the SaaS Platform and host AI models. This includes:

- Operating and securing the underlying cloud infrastructure used to host AI models and services
- Providing secure model hosting, access management, and orchestration capabilities
- Protecting customer data through built-in security, privacy, and resilience controls
- Delivering safety and responsible AI capabilities within applicable AWS services
- Maintaining physical security of AWS data centres
- Securing foundational cloud services, including networking, compute, storage, and database services
- Aligning AWS services with applicable security, privacy, and compliance required supported by AWS

AI Model Developer

TechnologyOne's AI enabled products leverage large language models (LLMs) and/or machine learning which are hosted on AWS. As the creator of these models, the model developer is responsible for designing, training and maintaining models in a manner that promotes safety, security, transparency, and reliability. This includes:

- Designing model architecture in line with responsible AI principles
- Curating and preparing training data to support appropriate and lawful model behaviour
- Training, evaluating, and testing models to validate performance and reliability
- Implementing model-level safety measures to reduce the risk of harmful or unintended outcomes
- Securing models against misuse, abuse, or unauthorised access
- Providing ongoing updates and maintenance throughout the model lifecycle
- Maintaining documentation and transparency regarding model capabilities, limitations, and intended use

Contact Us

If you require further information about this Policy, please contact the TechnologyOne RCS Council via AIGovernance@technology1.com

Review and Updates

This Policy was last updated on 1 July 2026 It will be reviewed and updated annually or when significant changes occur.

Further Information and Contacts

For more information or to provide feedback about TechnologyOne's products, services, and our approach to privacy and security, refer to the following resources:

- TechnologyOne website via www.technology1.com
- Privacy Policy via www.technology1.com/privacy-policy
- SaaS+ Security via www.technology1.com/saas-plus/security
- Chief Privacy Officer via privacy@technology1.com
- AI Governance team via AIGovernance@technology1.com

Glossary

- **Artificial Intelligence (AI)** - A machine-based system designed to operate with varying levels of autonomy that, for explicit or implicit objectives, infers from input data how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments
- **Machine Learning** - A set of techniques used within AI systems that enables models to learn patterns from data and improve their performance in generating outputs without being explicitly programmed for each task.
- **Agent** - An automated entity that perceives its environment, makes decisions, and takes actions to achieve defined objectives, often with varying levels of autonomy
- **Model Training** - The process of developing or improving an AI or machine learning model by using training data to determine and adjust the model's parameters so it can generate outputs such as predictions or decisions.
- **Model Developer** - A natural or legal person or entity that designs, develops, or has an AI system or model developed and places it on the market or into service under its name or control.