Legal

AI Policy

Effective 2026-09-21 / Updated 2026-09-21

This AI Policy describes how Trickest, Inc. ("**Trickest**", "**we**") approaches artificial intelligence in the Trickest platform. It is written for our customers and anyone evaluating Trickest. It is provided for information; see section 7 below.

In short

- We do not use Customer Content to train or fine-tune AI models. The settings we apply to every AI request we make route it only to model providers that have agreed with our AI gateway provider not to train on it.
- AI requests that we make go through an AI gateway to model providers that have agreed not to train on them and, apart from the limited exceptions we disclose, not to retain them.
- The Trickest agent works within the permissions of the user who runs it, in an isolated environment, and is designed to ask before destructive actions.
- You stay in control: you can ask us to restrict the models available in your workspace, use your own model provider key, or have AI turned off for your workspace.
- We do not use AI to make decisions about people.

1. Scope

This Policy covers the AI features of the Trickest platform, including the Trickest agent and the AI-assisted features that support it, such as summaries, workspace memory, and search. It does not cover Trickest's workflow engine, which runs security tools deterministically, except for AI features used in or with workflows, such as the AI Agent workflow node, which this Policy does cover. Capitalised terms not defined here have the meaning given in your agreement with Trickest.

In this Policy, "**Customer Content**" means the prompts, files, and other inputs you or your users provide to AI features, the content AI features retrieve from your systems or authorised targets while performing a task for you (such as tool output and scan results), and the outputs those features generate for you, including data derived from them such as embeddings and memory.

2. Our commitments

2.1 No training on your data

We do not use Customer Content to train or fine-tune AI models. The settings we apply to every AI request we make route it only to model providers that have agreed with our AI gateway provider not to train on it. We use Customer Content only to provide, support, and secure the platform for you and to comply with law: for example, to answer a support request you raise or to investigate a security issue. Access to Customer Content by Trickest personnel is limited to those who need it for these purposes.

2.2 Model providers

The AI features use large language models developed and hosted by third-party model providers. Unless you use your own model provider credentials, we access these models under our own accounts through an AI gateway, which routes each request to a model provider. Our sub-processors, the model providers the gateway may forward requests to, and the processing locations are listed on our sub-processor list.

We make a model available in the platform only where the model provider has agreed, with us or with our AI gateway provider, not to train on Customer Content and not to retain it beyond processing the request, subject to the exceptions below; where processing takes place in the European Economic Area, the United States, or another country with appropriate safeguards under the GDPR; and after we have evaluated the model for quality and security.

Some providers may keep content that their safety systems flag as abusive, for a limited period, to enforce their usage policies, and a provider may keep content where the law or legal process requires it. Web search and voice dictation use services that have agreed not to train on the content but do not offer zero data retention, so the provider may keep search queries or audio for a limited period under its terms.

If you use your own model provider credentials, requests made with those credentials go to that provider under your own agreement with it, and that provider is not our sub-processor. Other AI features, such as summaries, search, and dictation, may continue to use models that we access through our AI gateway.

2.3 Data location and retention

The platform stores Customer Content in Amazon Web Services in Frankfurt, Germany, unless your agreement specifies another region. Transfers outside the European Economic Area rely on the safeguards described in our Privacy Policy and data processing agreement. By default, we keep conversations for the life of your account and permanently delete removed chats within 30 days;

workspace memory logs are kept for 90 days, session memory and search indexes for 120 days, and operational logs, which exclude prompt and output text, for 15 days. Your data processing agreement, where you have one, sets out the retention periods and deletion commitments that bind us.

2.4 Security

The AI features operate within Trickest's information security management system, which is certified to ISO/IEC 27001. The security measures we are contractually bound to maintain are set out in our data processing agreement.

The Trickest agent is designed to run commands only in isolated environments dedicated to a single chat or workflow, with access limited to what the task requires and without access to model provider credentials or to Trickest's internal systems. Customer Content is logically separated by workspace. Agent actions are subject to permission controls, are logged, and are bounded by limits on duration, steps, usage, and spend.

Web pages, files, and scan results that the agent processes may contain instructions intended to manipulate it (prompt injection). The agent is designed to treat such content as data, not as instructions, and the controls above are intended to limit the effect of any manipulation. No AI system is fully resistant to prompt injection.

Further detail on our security architecture is available to customers and prospective customers under a non-disclosure agreement.

2.5 Human oversight and your controls

Every agent session is started by a user, or by a schedule or integration a user has set up. Users can see the agent's steps and stop a run in progress. By default, the agent is designed to ask a user before actions we classify as destructive, and to decline such actions where approval cannot be requested, for example in scheduled runs and in chats through Slack or Telegram. Workspace administrators can set further rules, through our API, for which actions are allowed, need approval, or are denied.

On request, we can restrict the models available in your workspace, turn off web search and dictation, or turn off AI for your workspace entirely. The rest of the platform remains available.

2.6 Transparency

We tell users when they are interacting with an AI system. We maintain an AI Transparency Note that explains how the agent works, what data AI models receive, and how long it is kept; it is available to customers on request.

3. Your responsibilities

When you use the AI features, you are responsible for:

- using the agent only against systems you own or are authorised to test, and in accordance with our Acceptable Use Policy;
- reviewing AI outputs, and the actions the agent proposes, before relying on them or acting on them;
- configuring permissions for your workspace and users appropriately;
- not submitting special categories of personal data, criminal-offence data, or records about your employees, payroll, or your own customers to AI features unless your agreement with us covers that processing; and
- if you use your own model provider credentials, complying with that provider's terms.

4. Limitations of AI

AI outputs can be inaccurate, incomplete, or inconsistent, and the same request can produce different results. Findings from the Trickest agent do not guarantee that a system is free of vulnerabilities. AI features are provided to assist security professionals, not to replace their judgement.

5. Governance

Trickest's leadership owns our approach to AI and reviews it at least once a year and when our services, providers, or applicable law change. New AI features and model providers are assessed for risk before release. Our AI governance is aligned with ISO/IEC 42001 and the NIST AI Risk Management Framework; Trickest is not certified to ISO/IEC 42001. Our staff receive training on the responsible use of AI.

Under the EU AI Act, Trickest is the provider of the Trickest agent, which is intended for security automation and is not used to make decisions about individuals. We assess our obligations under the Act, including its transparency requirements, as they apply.

6. Reporting concerns

If you believe the Trickest agent has been misused, or you have a concern about how our AI features behave, contact security@trickest.com. We may suspend AI features for a workspace where needed to protect the platform or others.

7. Relationship to your agreement

This Policy describes Trickest's approach to AI as at its effective date and is provided for information. It does not form part of, amend, or extend any agreement between Trickest and a customer, including the Master Software and Services Agreement, the Terms of Use, any AI terms addendum, and any data processing agreement, and it is not one of the "Policies" incorporated by reference into the Terms of Use. Our binding commitments are those set out in those agreements. If this Policy conflicts with them, the data processing agreement prevails as to the processing of personal data, and otherwise the agreement prevails. Nothing in this Policy is a warranty or guarantee of the performance, accuracy, or security of any AI feature.

8. Changes

We may update this Policy to reflect changes to the platform, our providers, or applicable law. We will publish the updated version with a new effective date and keep a record of changes. If we make a material change to how we use Customer Content, including for training, we will notify customers at least 30 days in advance, and the change will apply only to Customer Content provided after it takes effect. Changes to sub-processors are notified as set out in our data processing agreement.

9. Contact

Questions about this Policy or about Trickest's use of AI: privacy@trickest.com. Trickest, Inc., 1111B S Governors Ave STE 6511, Dover, DE 19904, United States.

Related documents: Terms of Use · Privacy Policy · Acceptable Use Policy · Sub-processors · Service Level Agreement