Rapport AI Use Statement

Purpose, data flows and controls for Rapport's AI Features

This statement describes how Rapport, a Service of Actionable Science, Inc. ("Rezolve.ai", "we"), uses large language models: the intended purpose, the architecture and data flows, the controls that limit what the AI can do, and the obligations of customers who enable it. It supplements the Rapport Product Terms (Section 8), the Privacy Policy and the Data Processing Addendum. Capitalised terms have the meanings given in the Rapport Product Terms.

1. Summary

Topic	Position
Default	All AI Features are off for new tenants until an Administrator enables them. One tenant setting controls all AI Features.
Provider	Anthropic (Claude models). No other AI provider is used for Rapport.
Provider retention and training	Under our agreement with Anthropic, inputs and outputs are not retained after processing and are not used to train models.
Our use of data	We do not use Customer Data, prompts or outputs to train or fine-tune any model.
Device changes	The AI can propose Device Actions but cannot make changes on its own. Modifying or elevated commands are held for human approval; destructive or disallowed commands are refused.
Automatic diagnostic access	Once AI Features are enabled, the AI Assistant may collect file metadata (not file contents) and run commands classified as read-only without separate approval, where permitted by the Customer's settings. Information retrieved may be sent to Anthropic as described below.
Records	Every prompt, tool call, result, approval and decision is recorded in the session transcript and the Audit Log.

Data location

Session transcripts are stored in the United States. Authorised support personnel in India may access relevant data as described in the Data Processing Addendum and Sub-processor List.

2. Intended purpose and prohibited uses

Rapport's AI Features are intended to help authorised IT staff diagnose and resolve technical issues on devices their organisation manages. They are not designed or intended to monitor, evaluate, profile or make decisions about individuals, including employees, and must not be used for those purposes. Customers must not use AI output as the sole basis for any decision with legal or similarly significant effects on an individual.

3. AI Features

Feature	Function	Data sent to the model	Model account
AI Assistant	Investigates a specific Enrolled Device by running Diagnostic Tools, explains findings and proposes Device Actions	Operator prompts; conversation context; tool definitions; device facts; Diagnostic Tool results	ASC's account
Investigations	Organises read-only diagnostic observations into a finding, a confidence level and next steps	Diagnostic observations for the device under investigation	ASC's account
Copilot	Answers Operator questions in the Rapport Console	Operator prompts and the console context needed to answer	ASC's account
Script drafting	Drafts scripts for Operator review	Operator instructions and script context	ASC's account

Feature	Function	Data sent to the model	Model account
Issue resolution	Suggests resolution steps for detected issues	Issue details and related device observations	ASC's account
Patch-risk assessment	Assesses the risk of a patch before deployment	Patch metadata and related device observations	ASC's account

4. Architecture and data flow

1. An Operator, authenticated with multi-factor authentication and scoped to the Customer's tenant, starts a session in the Rapport Console.
2. The Rapport gateway checks the tenant, the AI Feature setting, session limits and the Operator's permissions, then calls the Anthropic API using credentials held server-side. Provider credentials are never sent to devices or browsers.
3. For the AI Assistant, the model may request a Diagnostic Tool. The request is sent to the Agent on the selected Enrolled Device over its authenticated outbound HTTPS connection.
4. The Agent classifies the request locally, independently of the model (section 5), and runs, holds or refuses it.
5. The Agent returns bounded results to the gateway. Rapport masks text that matches credential patterns on a best-efforts basis before sending the results to the model for the next turn. Masking may not identify or remove all credentials or Personal Data.
6. Prompts, model responses, tool calls, results, approvals and decisions are stored in the session transcript in the Customer's tenant and recorded in the Audit Log.

5. Controls on device access

5.1 Command classification

Class	Examples	Handling
Read-only	Querying service state, reading configuration, listing processes	May run automatically where the Customer's settings permit
Modifying or elevated	Restarting a service, changing a setting, installing software	Held until an authorised Operator approves that specific command; approval does not extend to later commands
Destructive or disallowed	Commands matching Rapport's destructive or disallowed policy	Refused

Clipboard and screen capture

Reading the clipboard, taking a screenshot

Held for approval; if approved, the content passes through automated data loss prevention filtering and is sent to Anthropic

Classification is performed by the Agent on the device. A model's request alone is never permission to run a command.

5.2 File access

- Once AI Features are enabled, the AI Assistant can collect metadata about files the Agent can access (such as names, paths, sizes and dates) without separate approval where permitted by the Customer’s settings. It does not read the contents of files.
- Operating-system credential stores are always blocked.
- On devices running Rapport Agent 1.1 or later, the AI Assistant is also blocked from browser profile data and from email and messaging stores, including Apple Mail, Messages, Microsoft Outlook, Microsoft Teams and Slack, including through database query commands.
- Scripts and file transfers initiated by Operators are separate Device Actions. They are not subject to the AI Assistant's file block, require approval by default and are recorded in the Audit Log.

5.3 Limits

Administrators configure the model, conversation length and output budgets, and the Diagnostic Tools available to the AI Assistant. **Administrators can also limit the folders the AI Assistant can access.** The gateway enforces session limits and records usage.

6. Models

The AI Assistant uses Claude Opus by default. Administrators may select Claude Sonnet or Claude Haiku in Admin → AI settings. Other AI Features use the Claude model selected by Rezolve.ai for that feature. We may change the models used, within Anthropic's Claude family, to maintain quality and availability; a change of AI provider is subject to the sub-processor notice process in the Data Processing Addendum.

7. Provider terms

Anthropic acts as our sub-processor. Under our agreement with Anthropic, inputs sent through our account and the outputs generated are not retained by Anthropic after the response is returned and are not used to train Anthropic's models. All AI Features use our account; customer-supplied API keys are not supported.

8. Retention

AI Assistant session transcripts, including tool calls and results, are retained for 180 days by default. Administrators can set a period between 1 and 3,650 days. Audit Log entries are retained for at least 365 days. Customers must select retention periods appropriate to their purposes and applicable legal requirements. Deletion, including on device retirement or termination and the treatment of backups and legally required retention, is governed by the Data Processing Addendum.

9. Known limitations and risks

Risk	Mitigation
Inaccurate or incomplete findings	Findings show supporting evidence and a confidence level; Operators review before acting; a completed command is not treated as proof of resolution
Inappropriate proposed actions	Device-side classification; approval of each modifying command; refusal of destructive commands
Instructions hidden in device content (prompt injection)	Classification and approval run on the device independently of the model, so injected text cannot authorise a command
Personal data in device content	File blocks for credential, browser, email and messaging stores (Agent 1.1+); best-effort credential masking; customer control of AI enablement and tools
Credential leakage	Credential stores blocked; credential-pattern masking; provider keys held server-side

Read-only access can still expose confidential information or Personal Data. File restrictions, credential masking and approval controls reduce risk but do not guarantee that such information will never be accessed or disclosed.

10. Customer responsibilities

- Before enabling AI Features, ensure that the proposed use has a lawful basis, affected individuals receive the notices required by applicable law, and any required consents or consultations have been completed.
- Grant AI and approval permissions only to trusted Operators.
- Review findings and proposed Device Actions before acting.
- Use AI Features only for their intended purpose (section 2).
- Limit prompts, file access and diagnostic results to information necessary for the authorised support purpose, and comply with the restricted-data requirements in Section 5.7 of the Rapport Product Terms.

11. Contact

Questions about this statement: privacy@rezolve.ai. Security concerns about AI Features: security@rezolve.ai.