RIVAL
A RIVAL GROUP COMPANY

AUGUST 2026

THE RIVAL PLATFORM

Trustworthy AI for better research

AI is changing research. Trust determines whether it creates value.

01 Our approach to AI

02 How AI is used on the platform

03 Trusted third-party AI partners

04 AI guardrails

05 Enterprise security controls

06 Continuous improvement

SECTION 01

AI is changing research. Trust determines whether it creates value.

At Rival Technologies, our approach is simple. AI should make research better — faster to design, sharper to read, and more respectful of the people who take part in it. Every AI capability on the Rival Platform is built against four commitments.

01

AI should help people make better decisions

02

AI should increase transparency

03

AI should respect respondent privacy at every stage

04

AI should strengthen trust in research

> “At Rival, a human is involved in every design, every release, every innovation.”

SECTION 02

How is AI used on the Rival Platform?

There are two main methods for using AI today on the Rival Platform.

A Researcher tools

AI enhances the workflows researchers already use — survey design, brand and tone alignment, and gathering insights.

B Smart Probe

Smart Probe uses AI to interact with respondents in a controlled manner, deriving a deeper understanding of their preferences.

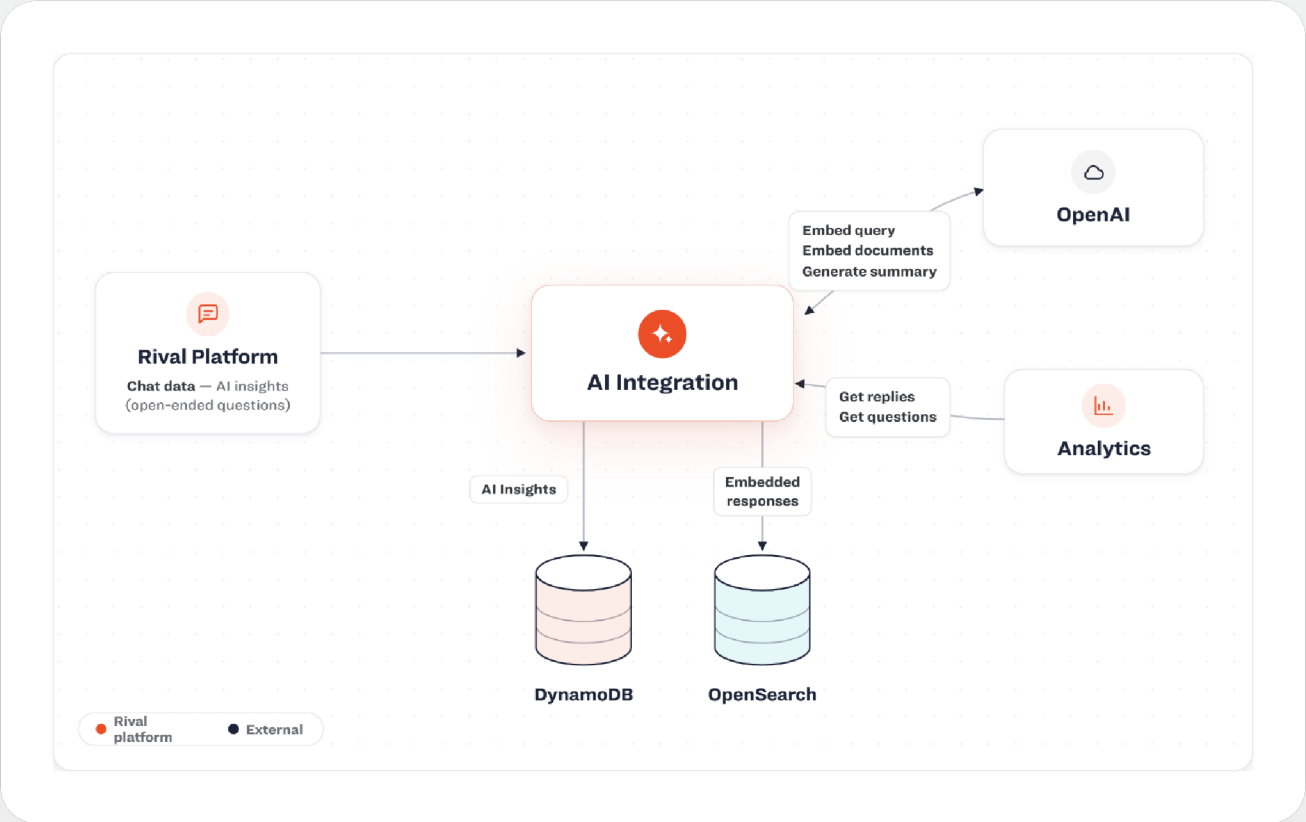
NON-NEGOTIABLE

Customer data is never used to train AI in any way, and our platform is designed to avoid sending PII to third-party AI providers.

FIGURE 01

Rival Platform AI features — data flow

Chat data from open-ended questions is the only participant content that reaches the AI integration layer. Embeddings and insights stay inside Rival infrastructure.



RIVAL PLATFORM

Chat data, AI insights, and embedded responses are stored and processed within Rival's own environment.

EXTERNAL

Third-party models receive scoped requests only — embed query, embed documents, generate summary.

SECTION 03

Trusted third-party AI partners

We only use trusted AI partners and tools for internal processes and features available on the Rival Platform. Current features like Smart Probe and Insight Summarizer leverage OpenAI's models and API.

All third-party AI providers must go through a thorough security and compliance review prior to use in our products and services.

MODEL CHOICE

When selecting models for our AI features, we put them through defined evaluation frameworks.

Quality
Does it hold up on real research data?

Accuracy
Does it represent what respondents said?

Safety
Does it stay inside our boundaries?

0 customer records used to train a third-party model — today, or in any future release.

SECTION 04

AI guardrails

Four layers of control sit between a research question and an AI response — each one owned, reviewed, and versioned.

01
Human-in-the-loop

02
Prompt protection

03
Input protection

04
Output protection

GUARDRAIL 01

Human-in-the-loop

Currently we empower humans through our features by:

01 **Tone Refinement** includes human-in-the-loop by design.

02 **The AI Summarizer** requires administrative review, and allows admins to edit or delete information before sharing.

03 **Smart Probe** operates within specific guardrails set up by a researcher when they design a survey.

GUARDRAIL 02

Prompt protection

Our AI services operate within a private, secured environment. All requests are authenticated and authorized before processing, and access to AI configuration — including system prompts — is restricted to authorized personnel only.

System prompts are centrally managed, version controlled, and designed with clear boundaries between system instructions and user-provided content. Privacy and security guidance is built into AI prompts to help ensure responses align with our data handling requirements.

GUARDRAIL 03

Input protection

AI features are designed to process participant open-ended survey responses and user-authored content only. Inputs are validated before processing, and prompt templates help ensure user content is incorporated safely and consistently.

We continue to invest in strengthening these controls as our AI capabilities grow.

To help protect service availability, AI requests are subject to rate limiting, usage monitoring, and resource controls. We also monitor AI usage patterns and maintain safeguards to prevent excessive or abusive use.

GUARDRAIL 04

Output protection

AI responses are validated before being returned to the application to help ensure they conform to expected formats and do not expose internal system information. Error messages are intentionally sanitized to prevent disclosure of technical details.

Content generated by AI that is published or shared externally is reviewed and approved by a human before use. For real-time AI capabilities, responses are limited to predefined topics and research objectives established during survey configuration, helping ensure interactions remain within intended boundaries.

SECTION 05

Enterprise security controls

Strong security practices are foundational to trustworthy research, and the use of AI doesn't change that.

Rival Technologies maintains a third-party certified security program designed to protect client data, respondent information, and platform operations. Audit reports and certifications are available upon request.

ISO 27001	SOC 2	HIPAA
:2022 certified	Type II audited	Controls audited annually

CONTROLS

Zero-trust architecture	All access requests are authenticated, authorized, and logged. No user or system is trusted by default — including AI agents.
Encryption	Data is encrypted both in transit and at rest using industry-standard security controls.
Data segregation	Segregation is meticulously maintained through a sophisticated system of user permissions controlled by our application code.
Controlled administrative access	Access to production systems is restricted, monitored, and governed through formal approval processes.
24/7 monitoring and alerting	Our SIEM provides continuous scanning and detection. Alerts for misconfigurations, vulnerabilities, or suspicious activity are routed to relevant personnel in real time.

SECTION 06

Continuous improvement, continuous innovation

We continue to explore new AI functionality and workflows to best serve our customers. Some AI capabilities may be offered as limited, opt-in pilots while they undergo additional evaluation, customer feedback, and security review before becoming generally available.

Pilot features on the Rival Platform are still subject to rigorous back-end controls and validation. Our teams adhere to an AI Acceptable Use policy, and undergo mandatory training on how to use AI safely — including targeted training for those in development roles.

Want the audit reports?

Our ISO 27001, SOC 2 Type II and HIPAA documentation is available on request. Talk to your Rival contact, or reach us at **security@rivaltech.com**.

© 2026 Rival Technologies. This document describes controls in place at the time of publication and may change as the platform evolves.

]]>