TriangulAIte

SECURITY OVERVIEW, EARLY ACCESS

Your work stays yours.

In early access, each firm's work runs in an environment of its own, with its own cloud account, database, file storage, key store, and sign-in. Nothing is shared with any other firm. Everything you send us is encrypted before it's stored, and every step of every analysis is written to a tamper-evident record. We build to SOC 2 and ISO/IEC 27001, and to NIST's frameworks for secure development and AI risk.

This describes early access, from late February 2027. The demo from mid-January runs on our own hardware, so bring documents that are no longer sensitive.

Your own environment

A dedicated stack for each firm. Each firm gets its own cloud account, database server, file storage, key store, sign-in service, and permission service. No table, database, server, storage bucket, or key is shared between firms.

Your stack stays in the United States. Early access runs in the US, and your database, files, backups, and keys stay there. Regions in the EU and the UK follow later.

Where it runs. In early access, your dedicated stack runs in our cloud. Running in your own cloud, or on a fully air-gapped workstation, comes after early access.

Encryption and keys

Encrypted before it's stored. Your documents, prompts, and results are encrypted with AES-256-GCM before they reach the database or file storage. A copy of the database alone reveals nothing.

Tamper-evident. Each encrypted value is bound to the record it belongs to, so a value that's altered, or moved to another record, won't open.

Your keys, in your key store. Keys are held in three layers in your firm's own key store. The root key sits in a cloud key service or a hardware security module. Each analysis has its own key, which never leaves the key store.

Encrypted in transit. Every connection inside the stack uses TLS, verified against your stack's own certificate authority.

No content in logs. Logs, traces, and monitoring carry identifiers and timings, never your content.

Deletion you can rely on. Deleting an analysis destroys its key, which makes its content unreadable everywhere at once, and backups age out within 30 days. Closing an account removes everything within 37 days, after a 7-day window to change your mind. A model provider may keep its own copy of what it was sent for a limited time, under its terms.

Access and identity

Sign in your way. Use your own identity provider, through single sign-on. A second factor is always required, either authenticator codes or passkeys, with no opt-out.

Every action is checked. Every operation checks who's asking, for which firm, and whether they're allowed, before it touches any data. Removing someone takes effect within a minute.

No stored passwords. Database passwords are created, held, and rotated by your key store. No part of the system keeps one.

Closed to our staff. Our staff see none of your content unless you grant access, for support or to build a custom workflow for you.

AI models

Nothing you upload trains a model. Analyses run on commercial AI models through their providers' paid APIs, under terms that don't train on what's sent.

You choose the providers. When your firm starts, it chooses which model providers may see its documents, and every call stays inside that list. Each step uses the best model from the providers you approve, and approving more gives a stronger analysis.

Every call is recorded. Each call to a model provider is recorded, with the provider, its region, and its retention terms, and goes only to an approved provider's own endpoint.

A record of every step

A record that only grows. Every analysis is written to a record that can be added to but never changed. The application may only add and read, and the database refuses any edit or deletion.

Chained and signed. Each entry is chained to the ones before it. Every minute, a checkpoint is signed with your firm's key, time-stamped by an independent authority, and stored for 7 years in a separate account, under a lock no one can lift.

Checked from outside. A daily check, run from outside the system, confirms the record matches its chain.

How we manage risk

Security is built into how we work, from the first design to every change.

01 **We find hazards before we build.** Every design records what could go wrong, such as exposed data, tampering, a misleading conclusion, or an AI failure, in one hazard log. Each entry records how severe it is, which parts it affects, and how it's mitigated. AI risks sit in the same log, following the NIST AI Risk Management Framework.

02 **We scale rigor to severity.** The parts whose failure would do the most harm are held to the strictest standard, with near-complete test coverage and property-based testing. Every change is reviewed by an AI reviewer and by an engineer before it merges.

03 **Every mitigation becomes a tested requirement.** Each mitigation is written as a requirement and linked to an automated test. The build fails if any requirement lacks a passing test, so the evidence that a control works is collected on every change.

04 **Our controls map to the standards.** Each safeguard is recorded against the NIST SP 800-53 catalog and against the SOC 2 and ISO/IEC 27001 requirements it answers. The same register is our ISO/IEC 27001 Statement of Applicability.

05 **Every change passes automated checks before it merges.** They cover secrets, known vulnerabilities, licenses, the software supply chain (pre-built packages only, and a 7-day wait on new releases), and code quality. Nothing merges with a failing check.

06 **We plan for incidents.** A written incident plan commits us to tell an affected firm within 72 hours, as the SEC's Regulation S-P requires of firms' service providers. For each firm, we keep a list of every provider we use, with its retention and region, to support the firm's own oversight.

Standards

BUILT TO	SOC 2, ISO/IEC 27001, NIST SP 800-53, the NIST Secure Software Development Framework (SP 800-218, with its generative AI profile), and the NIST AI Risk Management Framework
INFORMED BY	ICD 203 and 206, the intelligence community's standards for stating confidence and sourcing, and ISO/IEC 42001, on AI management

A SOC 2 Type II audit is planned. Our controls are built to run from the first early-access team's first day, so the audit's observation window can begin then.

One place to change. All cryptography goes through one module, so moving to validated cryptography for a customer that needs it changes one place in the system, with no redesign.

Questions from your compliance team? Write to steve@triangulaite.com, or book a call at triangulaite.com.