SECURELOGIC AI

SECURITY OVERVIEW

Version 1.0 | Effective Date: [INSERT DATE]

Last Updated: [INSERT DATE]

---

1. Purpose

This Security Overview describes the security practices, controls, and operational safeguards implemented by Threat Loom, LLC, doing business as SecureLogic AI ("SecureLogic AI," "we," "our," or "us").

This document is intended to provide customers, prospects, and stakeholders with a transparent understanding of how SecureLogic AI protects customer information and maintains secure operations. We have written it to be specific where we can credibly be specific and honest where we cannot — including about practices we are working to mature as the company grows.

This document is provided for informational purposes and does not constitute a contractual commitment unless expressly incorporated by reference into a written agreement between SecureLogic AI and a customer. It supplements, and does not replace, our Terms of Service, Privacy Policy, and AI Transparency and Responsible Use Policy.

2. Security Program

SecureLogic AI maintains a security program designed to support the confidentiality, integrity, and availability of information processed through the Services. The program is implemented through a combination of:

- Engineering and architectural controls embedded in the platform codebase;
- Operational practices governing how systems are configured, monitored, and changed;
- Third-party security infrastructure provided by our service providers and subprocessors;
- Periodic internal review of controls, vulnerabilities, and security-relevant events.

Security practices are reviewed and updated based on business requirements, emerging threats, industry best practices, customer expectations, regulatory developments, technology changes, and operational experience. This Security Overview reflects controls in place as of the Effective Date and will be updated as the program evolves.

3. Security Objectives

Our security program is organized around five core objectives:

- Confidentiality — protecting customer information from unauthorized access and disclosure;
- Integrity — preventing unauthorized modification of customer information and security-relevant records;
- Availability — maintaining reliable access to the Services for authorized users;
- Accountability — recording security-relevant events in immutable audit logs so actions can be reconstructed;
- Continuous improvement — identifying weaknesses through monitoring, review, and customer feedback, and acting on what we find.

4. Shared Responsibility Model

Security is a shared responsibility between SecureLogic AI and our customers.

4.1 SecureLogic AI is responsible for:

- Application and platform security, including secure coding practices and dependency management;
- Infrastructure configuration and hardening of cloud-hosted services;
- Authentication, authorization, and session management controls;
- Security monitoring, anomaly detection, and operator alerting;
- Vendor and subprocessor oversight;
- Incident response activities for incidents affecting our systems;
- Data protection safeguards for information processed through the Services.

4.2 Customers are responsible for:

- Securing endpoint devices used to access the Services;
- Managing user accounts within their organization, including provisioning, deprovisioning, and access reviews;
- Protecting credentials and enforcing strong authentication practices within their team;
- Enabling multi-factor authentication and, where appropriate, organization-level MFA requirements;
- Reviewing AI-generated outputs and Deliverables before relying on them, as further described in our AI Transparency and Responsible Use Policy;
- Avoiding submission of Prohibited Data (as defined in our Terms of Service);
- Maintaining their own compliance obligations and customer-side security controls.

5. Data Hosting and Processing Location

Customer information is stored and processed within the United States. The Services operate from data center regions in U.S. East (Virginia) and U.S. West (Oregon), supported by our infrastructure provider Render and our object storage provider Cloudflare R2.

We do not currently offer data residency outside the United States. Customers located outside the United States should review the international data transfer provisions of our Privacy Policy.

6. Security Architecture

SecureLogic AI is built on a modern cloud-native architecture designed for security, scalability, and operational resilience:

- A Next.js web application that serves the customer portal and manages authenticated user sessions;
- A separate API engine service that owns credential verification, business logic, and database access — separating browser-facing concerns from privileged operations;
- Background worker services that process asynchronous tasks (intelligence brief generation, posture scanning) without interrupting customer-facing performance;
- Managed PostgreSQL databases with encrypted connections, automated backups, and point-in-time recovery provided by our infrastructure provider;
- Cloudflare-fronted networking that provides DDoS protection, web application firewall capabilities, and a content delivery edge in front of our origin services;
- Cloudflare R2 object storage for customer-uploaded documents, with per-organization key prefixes that prevent cross-tenant access at the storage layer;
- Application monitoring via Sentry, with sensitive request data redacted before transmission.

Cross-tenant data isolation is enforced at the application layer through organization-scoped queries and row-level checks. We maintain a dedicated test suite ("cross-org-isolation") that runs on every pull request to detect any code change that could allow data from one customer to be visible to another.

7. Data Protection

7.1 Encryption in Transit

All connections between customers and the Services use Transport Layer Security (TLS). TLS termination is handled at the network edge by our infrastructure provider. Internal database connections between our application services and managed PostgreSQL also use TLS-encrypted connections.

7.2 Encryption at Rest

Data at rest is encrypted by our infrastructure providers. Managed PostgreSQL databases use Render-managed disk encryption. Object storage in Cloudflare R2 uses Cloudflare-managed encryption at rest.

7.3 Application-Layer Encryption of Sensitive Fields

Specific high-sensitivity fields within our databases are additionally encrypted at the application layer before being persisted, using AES-256-GCM authenticated encryption with a dedicated key separate from session signing keys. These include multi-factor authentication secrets and certain JSONB columns containing raw provider payloads, AI-generated report content, and intelligence brief content.

7.4 Cryptographic Hashing of Passwords

Customer passwords are never stored in plaintext or in any reversible form. We hash passwords using Argon2id with parameters that meet or exceed current OWASP recommendations (memory cost 64 MiB, time cost 3, parallelism 4). Argon2id is a memory-hard hashing function specifically designed to resist GPU-based brute-force attacks.

7.5 Secrets Management

All sensitive secrets — including API keys, database credentials, webhook signing secrets, and encryption keys — are sourced from infrastructure-provider environment variables and are never committed to source code or build artifacts. No credentials are hardcoded in the application codebase.

8. Authentication and Access Controls

8.1 Password Requirements

Customer passwords must be at least 12 characters long and include a mix of uppercase letters, lowercase letters, and digits. Newly set passwords are checked against a history of recent prior passwords to prevent cycling between a small set of credentials.

8.2 Multi-Factor Authentication

SecureLogic AI supports Multi-Factor Authentication (MFA) for customer accounts using time-based one-time passwords (TOTP) compatible with standard authenticator applications. Organization administrators may require MFA for all members of their organization. MFA secrets are encrypted at rest with AES-256-GCM authenticated encryption. Recovery codes are issued at the time of MFA enrollment, are hashed before storage, and are valid for single use only.

8.3 Account Lockout

Accounts that experience five failed login attempts are locked for a fifteen-minute window. Locked accounts can be unlocked by an authorized administrator before the lockout expires. Failed login attempts are recorded in the security audit log, and account lockout events trigger real-time operator alerts.

8.4 Session Management

Authenticated sessions use signed access tokens with a fixed maximum lifetime. Session tokens are invalidated when a user changes their password. Browser sessions are protected by cookies with appropriate security attributes (HttpOnly, Secure in production, SameSite restrictions).

8.5 Single Sign-On (SSO)

SecureLogic AI supports SAML-based single sign-on for organizations that wish to authenticate users through a corporate identity provider. SSO configuration is managed by organization administrators.

8.6 API Key Authentication

Programmatic access to the Services uses API keys that are hashed before storage; the plaintext key value is shown to the user exactly once at creation and is not recoverable from the platform afterward. API keys can be revoked at any time by their owner. API keys may optionally be issued with expiration dates.

8.7 Personnel Access

Access to production systems by SecureLogic AI personnel is restricted to authorized personnel with a legitimate business need and is subject to the same authentication requirements as customer access. Personnel access activities are logged and reviewable. Given our current organizational size, the number of personnel with production access is small; access controls will scale as the team grows.

9. Logging and Security Audit

9.1 Audit Log Coverage

Security-relevant events are recorded in a dedicated audit log. The audit log captures over 90 event types including authentication events (successful login, failed login, account lockout, password change, password reset), authorization events, API key creation and revocation, security configuration changes, anomaly detections, and administrative actions.

9.2 Audit Log Immutability

The security audit log table is protected by database triggers that prevent updates, deletions, and truncation by any database role, including administrative roles. This control is enforced at the database layer (not the application layer), so audit records cannot be tampered with even if application-level controls were bypassed.

9.3 Personal Information in Logs

Where event payloads would otherwise contain identifying information such as email addresses, the values are masked or partially redacted before being written to the audit log. Application error monitoring (Sentry) is configured with scrubbing rules that strip request bodies, cookies, authorization headers, and fields with sensitive-sounding names (passwords, tokens, API keys, session tokens, MFA codes) before transmission.

9.4 Log Retention

Security audit logs are retained for a period of at least twelve (12) months from the date of the event, as described in our Privacy Policy. Customers may request additional information about retention periods by contacting privacy@securelogicai.com.

10. Security Monitoring and Alerting

10.1 Anomaly Detection

The Services include automated detection routines that scan the security audit log on a scheduled basis for patterns indicative of attack activity. Current detections include credential stuffing (a single IP attempting to authenticate against a large number of distinct accounts within a short window) and API key probing (a single IP submitting a large number of invalid API keys). When detection thresholds are met, alerts are dispatched to operator channels.

10.2 Operator Alerting

Account lockouts, anomaly detections, and other security-relevant events trigger real-time alerts to operator notification channels via secure outbound webhook. Alert delivery failures do not prevent the underlying event from being recorded in the immutable audit log.

10.3 Application Error Monitoring

The Services use Sentry for application error monitoring. Error events are transmitted to Sentry only after passing through scrubbing rules that strip sensitive request data and redact fields with sensitive-sounding names. Sentry maintains its own SOC 2 Type II attestation.

11. Abuse Prevention

The Services include multiple layers of abuse prevention:

- Per-endpoint rate limits on sensitive routes including login, signup, password reset, AI query endpoints, and webhook delivery routes;
- A global request rate limit applied to all routes;
- Account-level lockout following repeated failed authentication attempts;
- Anomaly-detection scans that identify cross-account credential stuffing and API key probing patterns;
- Network-edge protection through Cloudflare, including DDoS mitigation and bot detection capabilities;
- File upload validation including content-type checks, magic-byte verification, per-file size limits, and per-organization storage quotas.

12. Application Security

12.1 Secure Coding Practices

The Services are built with the following secure coding practices applied throughout the codebase:

- Database queries use parameterized statements throughout. No string-concatenated SQL is used in the application;
- Output encoding follows the framework defaults of Next.js and React, which automatically escape interpolated content to prevent cross-site scripting (XSS);
- Content Security Policy headers restrict the sources from which scripts, styles, and other resources may be loaded by the application;
- Standard security response headers are applied to all responses, including HTTP Strict Transport Security (HSTS), X-Frame-Options, X-Content-Type-Options, Referrer-Policy, and Permissions-Policy.

12.2 Server-Side Request Forgery (SSRF) Defense

Customer-configurable outbound destinations (webhook URLs) are validated through multiple defensive layers before any HTTP request is dispatched: HTTPS scheme requirement, hostname denylist for cloud-provider metadata services, DNS resolution of all A and AAAA records with rejection of any non-public-unicast IP address, pinned outbound connection agents to defend against DNS-rebinding attacks, redirect rejection, and timeout enforcement.

12.3 Webhook Security

Outbound webhooks delivered to customers are signed using HMAC-SHA256 with a per-endpoint secret and a timestamp, allowing customers to verify webhook authenticity and prevent replay attacks. Inbound webhooks from supported providers (including Stripe and Resend) are verified using each provider's documented signature scheme.

12.4 Cross-Tenant Isolation

Each customer organization's data is isolated through application-layer enforcement of organization-scoped queries. A dedicated automated test suite runs on every pull request to detect any code change that could enable cross-tenant data access. This test suite is a required pre-merge check.

13. AI Processing Security

The Services use third-party artificial intelligence providers (currently Anthropic for large language model services and OpenAI for speech-to-text transcription) to deliver AI-assisted functionality. AI processing is governed by additional safeguards:

- Customer Content is not used to train SecureLogic AI's proprietary models;
- Customer Content is not intentionally submitted to AI providers for foundation model training;

- AI provider connections use the same TLS protection and key management as other outbound integrations;
- Sentry scrubbing applies to AI provider integrations so that secrets and sensitive request fields are not transmitted to error monitoring;
- AI provider arrangements and the categories of data transmitted to each are documented in our AI Transparency and Responsible Use Policy.

Additional information about AI features, providers, customer rights, and AI-specific incident reporting is available in our AI Transparency and Responsible Use Policy.

14. Software Supply Chain Security

We apply the following controls to manage software supply chain risk:

- Dependency vulnerabilities are scanned automatically on every pull request via npm audit, with a policy of failing pull requests that introduce high or critical severity vulnerabilities;
- Dependabot is enabled across our primary repositories to surface dependency updates on a weekly basis;
- Lockfiles are committed to source control to ensure reproducible installations across environments;
- Continuous integration workflows pin third-party GitHub Actions to specific commit hashes rather than mutable tags;
- Six required pre-merge checks must pass before code can be merged to protected branches: typecheck, lint, test, audit, build, and cross-organization isolation.

15. Change Management

Material changes to the application codebase and infrastructure are managed through a structured workflow:

- All code changes are proposed as pull requests against protected branches;
- Required pre-merge checks (see Section 14) must pass before merge;
- Protected branches (main, develop) enforce non-fast-forward merges and prevent branch deletion;
- Production deployments are gated by successful merge to the main branch, which triggers an automated build, migration, and deployment sequence;
- Database schema changes are managed through versioned migration files that are applied automatically on each deployment, with all migration outcomes recorded.

16. Vendor and Subprocessor Security

We rely on a limited set of third-party service providers to operate the Services. Our current subprocessors and a summary of the services each provides are listed in our Privacy Policy. The primary subprocessors include:

- Anthropic, PBC — large language model AI services;
- OpenAI, OpenAI Global, LLC — speech-to-text transcription services;
- Stripe, Inc. — payment processing (Stripe maintains PCI DSS Level 1 compliance);
- Render Services, Inc. — cloud application hosting and managed databases;
- Cloudflare, Inc. — content delivery, object storage, and edge security;
- Resend, Inc. — transactional and marketing email delivery;
- Functional Software, Inc. d/b/a Sentry — application error monitoring.

Each subprocessor is selected based on factors including their published security posture, applicable compliance attestations, and contractual data handling commitments. We update our subprocessor list as relationships evolve.

17. Availability and Resilience

Service availability is supported by our infrastructure providers and by operational practices within the Services themselves:

- Managed PostgreSQL databases are provided by Render, which performs automated backups and provides point-in-time recovery capabilities;
- Cloud-hosted services can be scaled to additional replicas as load requires;
- The platform is deployed across multiple regions (currently U.S. East and U.S. West) to support availability;
- Health checks are exposed by the application services so that infrastructure providers can route around unhealthy instances;
- Migrations and code changes are versioned and reversible; deployments can be rolled back if a regression is detected.

Specific recovery time objective (RTO) and recovery point objective (RPO) targets are not currently published; documented RTO/RPO targets are part of our planned operational maturity as the platform scales.

18. Backup and Data Recovery

Production PostgreSQL databases are backed up automatically by our infrastructure provider, Render. Render's managed database service supports point-in-time recovery within a defined retention window.

Object storage in Cloudflare R2 is durable by design through provider-managed redundancy. Customer-uploaded files are stored with per-organization key prefixes that prevent cross-tenant access at the storage layer.

Customers seeking specific backup retention windows or recovery SLAs for inclusion in a contractual agreement should contact security@securelogicai.com.

19. Incident Response

19.1 Detection

Security-relevant events are detected through a combination of application-layer logging (security audit log), application monitoring (Sentry), automated anomaly-detection scans, and operator alerting via webhook delivery.

19.2 Response Activities

When SecureLogic AI becomes aware of a suspected or confirmed security incident, response activities may include: evaluating the nature of the incident and systems affected, identifying the categories of information potentially involved, containing the incident, remediating affected systems, recovering normal operations, preserving evidence, conducting post-incident review, and notifying affected parties as described in Section 19.3.

19.3 Customer Notification

Where required by applicable law, contractual obligations, or our reasonable assessment of customer impact, SecureLogic AI will notify affected customers of confirmed security incidents within a commercially reasonable timeframe and provide such information as is reasonably available about the nature of the incident, the data potentially involved, and steps customers may take in response.

19.4 Formalization

Documented incident response procedures and tabletop exercises are part of our planned operational maturity. As the company grows, the incident response process will be formalized into a written runbook with documented escalation paths and stakeholder communication templates.

20. Vulnerability Disclosure and Security Reporting

SecureLogic AI welcomes responsible disclosure of suspected security vulnerabilities by security researchers, customers, and members of the public.

20.1 How to Report

Suspected vulnerabilities or security concerns may be reported by email to security@securelogicai.com. Please include sufficient detail to allow us to reproduce the issue, including any affected URLs, request payloads, and reproduction steps.

20.2 Our Commitments

When a vulnerability is reported in good faith and in accordance with this policy:

- We will acknowledge receipt of the report within five (5) business days;
- We will investigate the report and provide a substantive response within thirty (30) days, subject to reasonable extension where investigation requires additional time;
- We will not pursue legal action against researchers who comply with this policy and act in good faith, including by giving us a reasonable opportunity to remediate before public disclosure, refraining from accessing or modifying data that does not belong to them, and not degrading the availability of the Services for other users;
- Where appropriate, and with the researcher's permission, we may credit the researcher in our communications about a remediated vulnerability.

20.3 Bug Bounty Program

SecureLogic AI does not currently operate a formal bug bounty program with monetary rewards. A bug bounty program may be established as the platform matures.

21. Customer Data Rights and Control

Customers retain ownership of Customer Content submitted to the Services. Detailed customer rights — including rights to access, correct, delete, and obtain copies of personal information — are described in our Privacy Policy.

In addition, customers may:

- Manage their organization's users and roles directly through the customer portal;
- Enable or require Multi-Factor Authentication at the organization level;
- Configure SAML-based single sign-on for their organization;
- Generate, view, and revoke API keys for programmatic access;
- Request export of organizational data or termination of their account by contacting privacy@securelogicai.com.

22. Compliance Context

SecureLogic AI does not currently hold independent compliance certifications such as SOC 2, ISO 27001, HIPAA, PCI DSS, or FedRAMP authorization. As an early-stage company building toward those certifications, we rely on a combination of:

- The compliance posture of our underlying infrastructure providers — including Render, Cloudflare, Stripe, Anthropic, OpenAI, Sentry, and Resend — many of whom maintain SOC 2 Type II, ISO 27001, or equivalent attestations;
- The engineering and operational controls described throughout this Security Overview, several of which are designed to align with common control families (NIST SP 800-53, ISO 27001 Annex A) even where formal certification is not yet pursued;
- Transparent disclosure of our current posture, including practices that are not yet in place.

Our subprocessors' compliance attestations may be referenced as part of customer due diligence. Customers seeking specific documentation regarding subprocessor compliance posture may contact security@securelogicai.com.

We acknowledge that the absence of independent certification is a meaningful gap for some customers and use cases. We are committed to maturing our compliance posture as the platform and customer base grow.

23. Prohibited Data

Consistent with our Terms of Service and Privacy Policy, the Services are not intended or configured to receive, process, store, or analyze certain categories of restricted data — including Protected Health Information (PHI) regulated under HIPAA, Social Security Numbers, government-issued identifiers, biometric data, full payment card numbers (other than via our payment processor), student education records subject to FERPA, classified information, or special-category personal data under GDPR Article 9.

Customers are responsible for ensuring that Customer Content submitted to the Services does not include Prohibited Data. Submission of Prohibited Data violates our Terms of Service. If you become aware that Prohibited Data has been submitted, please contact security@securelogicai.com immediately so we can work with you to remediate.

24. Security Maturity and Roadmap

In the interest of transparency, we identify below several practices that are not yet in place but that we plan to mature as the company and customer base grow. We share this roadmap so that customers can calibrate their expectations and so that we can be held accountable to our stated commitments:

- Independent third-party penetration testing — currently not performed; planned as the platform approaches general availability;

- SOC 2 Type II attestation — not currently held; a planned milestone as enterprise customer demand requires it;
- Documented Incident Response runbook with stakeholder communication templates — currently informal; planned formalization within operational maturity work;
- Documented Business Continuity and Disaster Recovery plans with published RTO/RPO targets — currently relies on infrastructure-provider capabilities;
- Formal bug bounty program with monetary rewards — currently not operated; vulnerability disclosure handled via security@securelogicai.com (see Section 20);
- Bug-bounty-grade external pen testing cadence — not currently scheduled.

We update this section as items are completed or as the planned roadmap evolves. Customers with specific compliance or assurance needs are encouraged to discuss those needs with us directly.

25. Scope and Limitations

This Security Overview is intended to provide a high-level summary of security practices and intentionally omits operational specifics that would reduce the effectiveness of those practices if publicly disclosed (for example, specific detection thresholds, allowlist contents, internal architectural diagrams, and individual personnel responsibilities).

Customers with specific assurance needs that require deeper documentation may request additional information by contacting security@securelogicai.com. We may provide additional documentation under appropriate confidentiality terms.

This document does not create any contractual commitment unless expressly incorporated by reference into a written agreement between SecureLogic AI and a customer. Security practices are subject to change as the program evolves; the version date at the top of this document indicates the version under which the document was published.

26. Contact Information

For security inquiries, vulnerability disclosure, or to request additional security documentation:

Threat Loom, LLC

Doing business as: SecureLogic AI

44 Apple Street, First Floor

Tinton Falls, New Jersey 07724

United States

Security inquiries and vulnerability disclosure: security@securelogicai.com

Privacy inquiries: privacy@securelogicai.com

AI governance inquiries: ai@securelogicai.com

Legal and general inquiries: legal@securelogicai.com

27. Effective Date and Updates

This Security Overview is effective as of the Effective Date identified at the top of this document. We may update this Security Overview from time to time to reflect changes in our practices, infrastructure, or compliance posture. Material changes will be indicated by an updated Effective Date at the top of the document. Prior versions are archived and available upon request to security@securelogicai.com.

— End of Security Overview —