Strand Intelligence

AI Incident Response

Security & Data Protection Overview

Infrastructure Security · Data Residency · Access Controls

Version	1.0
Date	4 March 2026
Classification	Confidential – Approved for External Distribution
Prepared by	Strand Intelligence Ltd
Contact	will@strandintelligence.com
Web	strandintelligence.com

1. Executive Summary

Strand Intelligence is a digital forensics and incident response (DFIR) platform that enables security teams to conduct AI-powered investigations across endpoint and cloud telemetry. Given the sensitivity of forensic evidence and investigation findings, security is foundational to every layer of our architecture.

This document provides an overview of the security controls, data residency commitments, access management policies, and encryption standards that protect data within the Strand Intelligence platform. It is intended for distribution to external stakeholders who require assurance that data processed by Strand is handled with the highest level of care.

Platform	strandintelligence.com
Data Residency	United Kingdom (AWS London, eu-west-2)
Encryption at Rest	AES-256
Encryption in Transit	TLS 1.2+
Backend Access Model	IP Allowlisting + mTLS + MFA
Database Access Control	Row-Level Security with JWT Verification
WAF Layers	Dual Web Application Firewalls (no direct-to-origin)

2. Architecture Overview

The Strand Intelligence platform is composed of distinct infrastructure tiers, each purpose-built for its function and secured independently. This defence-in-depth approach ensures that compromise of any single tier does not grant access to another.

2.1 Frontend (Web Application)

The user-facing web application is a Next.js application hosted on Vercel. Vercel distributes static assets and serverless functions across its global edge network, meaning that application code executes as close to the end user as possible to minimise latency. Importantly, no customer data is stored or cached at the edge. All data requests are routed back to our UK-based backend infrastructure, with responses served over TLS-encrypted connections.

The frontend is responsible for presenting investigation findings, managing user sessions, and providing the investigation interface. Authentication is handled via Supabase Auth, which issues cryptographically signed JWTs upon successful login.

2.2 Investigation Backend (Evidence Processing)

The core of Strand Intelligence is our agentic AI investigation engine. This system runs on dedicated virtual private servers provisioned within AWS and DigitalOcean, all located in UK data centres. These servers are responsible for querying ingested forensic telemetry from our Clickhouse analytical database, performing automated investigation workflows, and generating investigation findings.

This backend infrastructure is entirely segregated from the public internet and is not accessible via any publicly routable address. Access is restricted exclusively to the Strand office IP address via allowlisting, with additional layers of authentication and encryption described in Section 4.

2.3 Data Stores

Clickhouse (Analytical Database): Stores ingested forensic telemetry, including endpoint logs, Microsoft 365 audit data, and other evidence sources. Clickhouse is optimised for high-throughput analytical queries over large datasets, making it ideal for forensic timeline reconstruction and pattern detection. Hosted within AWS London (eu-west-2).

PostgreSQL via Supabase (Application Database): Stores investigation findings, user accounts, organisation configuration, and access control metadata. This is the only data store directly accessible to the frontend web application, and all access is mediated through row-level security policies (see Section 5).

3. Data Residency

All customer data, forensic evidence, investigation findings, and supporting infrastructure resides within the United Kingdom. Specifically, all databases (PostgreSQL and Clickhouse), virtual machines used for evidence processing, and storage volumes are hosted in AWS London (eu-west-2) data centres.

The sole exception is the frontend application code distributed via Vercel’s edge network. This code contains no customer data and serves only to reduce page-load latency for users in different geographies. All data retrieval from the frontend is performed via API calls routed to our UK-based backend, ensuring that customer data never leaves UK infrastructure.

For further detail on how Vercel's CDN operates, refer to their public documentation at vercel.com/docs/cdn.

4. Backend Access Controls

Access to the investigation backend and underlying data stores is protected by multiple independent security layers, each of which must be satisfied before a connection is established. This layered approach ensures that no single point of failure can result in unauthorised access.

4.1 Network-Level Controls

IP Allowlisting: All inbound connections to backend infrastructure are restricted to the Strand office IP address. Traffic from any other source is dropped at the network boundary before reaching any application layer.

Dual Web Application Firewall (WAF): Two web application firewalls are deployed in series. The first WAF inspects and filters inbound traffic to the backend infrastructure. The second WAF is configured to drop any connection that does not originate from the first WAF. This architecture eliminates the possibility of direct-to-origin connections over the public internet, providing defence against origin IP disclosure and bypass attacks.

Blackhole Routing: Each server and database operates a blackhole configuration in which traffic that does not match a named, known internal route is silently dropped and an alert is generated. This prevents reconnaissance and reduces the attack surface by ensuring that only explicitly defined endpoints are reachable.

4.2 Authentication Requirements

Even from within the allowlisted network, authentication to backend infrastructure requires all of the following:

1. **Client Certificate (mTLS):** A client-side TLS certificate (mutual TLS) stored on authorised development devices and backed up to local, offline hardware. This certificate must be presented during the TLS handshake before any application-layer communication occurs.
2. **Username and Password:** Standard credential-based authentication with strong password policies.
3. **OTP-Based Multi-Factor Authentication:** Time-based one-time password (TOTP) multi-factor authentication, requiring possession of a registered mobile device.

The combination of these controls means that to gain unauthorised access, a threat actor would need to simultaneously be present on the Strand office network, in possession of an authorised development device with the correct client certificate, and have compromised both credential stores (laptop and mobile phone). Each additional factor exponentially increases the difficulty of a successful attack.

5. Database Security & Access Control

5.1 Encryption

- **Encryption at Rest:** AES-256 encryption at rest for all data stored in PostgreSQL.
- **Encryption in Transit:** All connections to databases are encrypted using TLS 1.2 or higher, ensuring data integrity and confidentiality in transit.

5.2 Row-Level Security (RLS)

The PostgreSQL database that serves investigation findings to the frontend application enforces row-level security (RLS) on every query. This is a database-engine-level control, not an application-level filter, meaning it cannot be bypassed by application vulnerabilities.

Before any data is returned, the following checks are evaluated by the database engine:

- **User Authentication:** The requesting user must be authenticated and hold a valid session.
- **Organisation Membership:** The user must belong to the organisation that owns the requested data.
- **Investigation-Level Access:** The user must have been explicitly granted access to the specific investigation.
- **JWT Verification:** The request must include a cryptographically signed JSON Web Token (JWT), verified against Supabase Auth's signing keys. Tokens are short-lived and cannot be forged without access to the signing secret.

If any of these checks fail, the query returns zero rows. There is no error message or indication of what data exists, preventing information leakage through failed access attempts.

6. Frontend Security

The frontend web application is secured through the following measures:

- **Authentication:** User sessions are managed via Supabase Auth, which handles credential storage, password hashing (bcrypt), and JWT issuance. No plaintext passwords are stored.
- **Transport Security:** All communication between the frontend and backend services occurs over HTTPS with TLS 1.2+. HTTP Strict Transport Security (HSTS) headers are enforced.
- **Hosting Security:** Vercel's hosting platform provides automatic DDoS mitigation, managed SSL certificates, and isolated serverless execution environments.

7. Operational Security

Beyond the technical controls described above, Strand Intelligence maintains the following operational security practices:

- **Least Privilege:** Access to production systems is limited to a minimal set of authorised personnel, following the principle of least privilege.
- **Offline Key Storage:** Client certificate backups are maintained on offline hardware, ensuring recoverability without exposing key material to network-based threats.
- **Monitoring & Alerting:** Blackhole routing alerts, WAF logs, and access logs are monitored for anomalous activity.
- **Infrastructure Segmentation:** Backend infrastructure is segmented such that the frontend application has no direct access to the investigation processing environment or the Clickhouse analytical database.

8. Summary

Strand Intelligence has been architected with security as a first-class concern at every layer. From network-level isolation and dual WAFs to mutual TLS authentication and database-engine-enforced row-level security, each control is designed to operate independently so that no single point of failure can compromise the confidentiality or integrity of client data.

All customer data resides within the United Kingdom, encrypted at rest with AES-256 and in transit via TLS. Access to backend systems requires physical presence on the office network, possession of a client certificate on an authorised device, valid credentials, and a one-time password from a registered mobile device.

For any further questions regarding our security posture, or to request additional documentation such as penetration test results or compliance certifications, please contact the Strand Intelligence team at will@strandintelligence.com.

Strand Intelligence Ltd | will@strandintelligence.com | 4 March 2026