# StriveDB Security Statement

Version 1.2 [Last Updated: March 2026]

## Why This Document Exists

StriveDB stores some of the most sensitive data in the nonprofit sector: survivor names, addresses, case notes, counseling records, forensic exam documentation, perpetrator information, and protective order details. A security failure in this context is not merely a compliance problem — it can endanger someone's physical safety. That reality shapes every architectural decision described in this document.

We believe transparency builds trust. This white paper describes what we do, how we do it, and where we are still strengthening our posture. We do not make claims we cannot substantiate.

# 1 Independent Security Assessment

StriveDB engaged **Accorian** (Esha IT Corp., East Brunswick, NJ) to conduct a gray box application penetration test in May 2025, with a full retest completed in August 2025.

The assessment included automated scanning, manual testing, and attempted exploitation across authentication, authorization, session management, input handling, and file upload functionality. Testing was conducted against the OWASP Top 10 (2021 edition) and additional threat categories, using multiple user roles (Admin, Case Manager, Counseling Director, Counselor) to evaluate privilege boundaries.

**Initial assessment results (May 2025):** 5 Medium, 2 Low, and 2 Informational findings were identified. No Critical or High severity vulnerabilities were found. The application was found to be safe against Broken Access Control, Cryptographic Failures, Security Misconfiguration, Vulnerable and Outdated Components, Server-Side Request Forgery, Software and Data Integrity Failures, and Security Logging and Monitoring Failures.

**Retest results (August 2025):** 7 of 9 findings were fully remediated. The remaining 2 low impact findings (rate limiting on secondary authentication endpoints, and edge-case file upload validation) were partially remediated at the time of retesting. **Both have since been fully resolved.**

Accorian rated StriveDB's overall web application security posture as **"Good"** (between Good and Excellent on their four-tier scale: Poor, Fair, Good, Excellent).

Key strengths noted by the assessors included encrypted communications, properly configured session cookies (HttpOnly and Secure flags), and well-configured security headers preventing clickjacking.

# 2 Infrastructure & Hosting

## 2.1 Cloud Environment

StriveDB is hosted entirely on **Amazon Web Services (AWS)** in the **US-West-2 (Oregon)** region. The infrastructure uses managed AWS services designed for security-sensitive workloads.

| **Component** | **AWS Service** | **Purpose** |
|---|---|---|
| Application | ECS<br>(Elastic Container Service) | Containerized application hosting |
| Database | RDS (PostgreSQL 16.8) | Survivor data storage |
| File Storage | S3 | Document and attachment storage |
| Container Registry | ECR | Docker image storage |
| Notifications | SNS | SMS-based MFA delivery |
| Threat Detection | GuardDuty | Continuous security monitoring |

## 2.2 Edge-Level Protection

Before any request reaches the StriveDB application, it passes through Cloudflare's security layer. Cloudflare provides **DDoS mitigation** that automatically absorbs volumetric attacks without impacting application availability, a **web application firewall** with managed rulesets that block common attack patterns including SQL injection, cross-site scripting, and remote code execution attempts at the network edge, **bot management** that distinguishes legitimate users from automated scanners and attack tools, and **rate limiting** at the edge that complements the application-level rate limiting described below.

This creates a **two-tier defense**: Cloudflare blocks known attack patterns and volumetric threats at the edge, while Rack::Attack (described below) handles application-specific rate limiting and behavioral detection closer to the application logic.

## 2.3 Network Isolation

The production database resides in a **private subnet** within an AWS Virtual Private Cloud (VPC) and is **not publicly accessible**. Database connections are accepted only from the application layer running in ECS. Direct infrastructure access (AWS console, deployment tools, database administration) requires authentication through an **OpenVPN** connection. All AWS IAM policies follow the **principle of least privilege**, granting only the minimum permissions required for each role and service.

## 2.4 Encryption

**Data at rest:** All survivor data stored in PostgreSQL is encrypted using **AES-256 encryption** via AWS RDS encryption. This covers the database instance, automated backups, snapshots, read replicas, and Multi-AZ failover replicas. The encryption is transparent to the application and cannot be disabled once enabled. File uploads stored in Amazon S3 are protected by server-side encryption. Authentication secrets (OAuth tokens, calendar integration credentials) are additionally encrypted at the application layer using **Rails Active Record Encryption (AES-256-GCM)**.

**Data in transit:** All connections between users and StriveDB are encrypted via **TLS**. HSTS (HTTP Strict Transport Security) is enforced at the load balancer level, instructing browsers to always use HTTPS. Internal connections between application containers and the database also use encrypted channels. Email-based MFA codes are transmitted via SMTP with STARTTLS encryption.

**Credentials management:** Application secrets (database credentials, API keys, encryption keys) are stored in **Rails encrypted credentials files** using environment-specific encryption keys. Production encryption keys are stored in a dedicated, access-restricted S3 bucket and are never committed to source control.

## 2.5 High Availability & Disaster Recovery

The production database is configured for Multi-AZ deployment, maintaining a synchronous standby replica in a separate AWS Availability Zone. In the event of an infrastructure failure, AWS automatically fails over to the standby with minimal interruption.

Automated daily backups of the database are enabled through RDS, with backups themselves encrypted at rest. These backups enable point-in-time recovery to any second within the retention window.

The S3 file storage service provides 99.999999999% (eleven nines) durability, automatically replicating objects across multiple facilities within the region.

## 2.6 Monitoring & Logging

- **AWS CloudTrail** is enabled across all regions, recording all API calls made to the AWS account for security auditing and forensic analysis
- **Amazon CloudWatch** provides real-time monitoring of application and infrastructure metrics, with alerting for anomalous conditions
- **Amazon GuardDuty** provides continuous, ML-powered threat detection — described in detail in the next section
- Application-level rate limiting and attack detection (described below) logs all blocked requests and suspicious patterns

## 2.7 Continuous Threat Detection

StriveDB uses **Amazon GuardDuty** for continuous, automated threat detection across the entire AWS environment. GuardDuty operates independently of the application — it monitors infrastructure-level activity that application logging cannot see.

**What GuardDuty monitors:**

GuardDuty continuously analyzes multiple data sources including AWS CloudTrail management and data events (detecting unauthorized API calls, credential misuse, and unusual account activity), VPC Flow Logs (detecting port scanning, communication with known malicious IPs, and anomalous network traffic patterns), and DNS query logs (detecting communication with command-and-control servers or cryptocurrency mining endpoints).

**How detection works:**

GuardDuty uses machine learning to establish a baseline of normal activity for the AWS account and then identifies deviations. It cross-references activity against continuously updated threat intelligence feeds maintained by AWS, including known malicious IP addresses, compromised domains, and attack signatures. This means new threat indicators are incorporated automatically without manual intervention.

**Examples of threats GuardDuty detects:**

- An attacker using stolen AWS credentials to access the environment from an unusual location or IP
- Attempts to escalate IAM privileges or disable security controls like CloudTrail logging
- An EC2 instance communicating with a known command-and-control server, suggesting compromise
- Brute-force attacks against SSH or RDP services
- Unauthorized attempts to access S3 buckets or modify bucket policies
- Reconnaissance activity such as unusual API call patterns that indicate an attacker mapping the environment
- Cryptocurrency mining on compromised compute resources

**Integration with incident response:**

GuardDuty findings are categorized by severity (Low, Medium, High) and feed directly into StriveDB's incident response procedures. High-severity findings trigger immediate investigation per the Incident Response Plan.

# 3 Authentication & Access Control

## 3.1 User Authentication

StriveDB uses a defense-in-depth approach to authentication with multiple reinforcing controls:

**No self-registration**. Users cannot create their own accounts. Every user must be invited by an organization administrator, ensuring that only authorized personnel gain access.

**Password requirements.** Passwords must be at least 8 characters and meet 3 of 4 complexity requirements (uppercase, lowercase, digit, special character). Passwords are hashed using bcrypt with a cost factor of 12, making brute-force attacks computationally impractical.

**Mandatory multi-factor authentication.** Every StriveDB user account requires at least one MFA method — this is not optional and cannot be disabled. Three MFA methods are supported: authenticator apps (TOTP, per RFC 6238), email-based one-time codes, and SMS-based one-time codes via AWS SNS. MFA codes are 6 digits, valid for 10 minutes, and invalidated immediately after successful use. A 30-second cooldown between resend requests prevents abuse.

**Account lockout.** After 5 consecutive failed login attempts, the account is automatically locked. Accounts can be unlocked via email link or automatically after 1 hour. The system warns users on their 4th failed attempt.

**Session management.** Sessions automatically expire after 4 hours of inactivity. All session cookies are flagged as Secure (HTTPS-only), HttpOnly (inaccessible to JavaScript), and SameSite=Lax (cross-site request protection). Logging out invalidates all session tokens, including "remember this device" tokens.

**Privacy-preserving error messages.** The system uses paranoid mode for authentication — login failures return a generic message that does not reveal whether the email address exists in the system.

**Login monitoring.** Every login is recorded with the IP address, browser, operating system, and geographic location. When a login occurs from an unrecognized device or location, the user receives an automatic email notification.

## 3.2 Role-Based Access Control

StriveDB uses a fail-secure, deny-by-default authorization system. Every request to access data or perform an action is evaluated against a policy. If no policy explicitly grants access, the request is denied. There are 41 separate authorization policies covering all data types in the system.

Access is governed by organization-defined roles with granular permissions:

Module access controls which areas of the system a user can enter (case management, counseling, volunteer management, reports). Victim access levels control which survivor records a user can see, ranging from no access to records they are personally assigned to, up to full access to all records. Service code permissions control which specific service types a user can log. Workflow permissions control actions like timesheet submission and approval.

Default roles are pre-configured for common positions (Administrators, Case Managers, Counselors, Counseling Directors, Volunteers), but organizations can customize role permissions to match their specific organizational structure and confidentiality requirements.

## 3.3 Survivor Record Restriction

Beyond role-based access, individual survivor records can be marked as restricted — visible only to specifically assigned providers and administrators. This supports high-risk cases where even standard role-based access is too broad, such as when a survivor is known to the organization's staff personally, or when the case involves a public figure or law enforcement officer.

# 4 Multi-Tenancy & Data Isolation

Each StriveDB customer organization operates in a completely isolated environment. No organization can see, search, access, or report on another organization's data under any circumstances.

## 4.1 How Isolation Works

Every database record is tagged with a tenant identifier. A system-wide security layer automatically adds a tenant filter to every database query — there is no way for application code to accidentally retrieve another organization's records. This isolation is enforced at the deepest level of the application: it applies to database queries, background jobs, file access, caching, and reporting.

Additional safeguards include: tenant identifiers are immutable after record creation (cannot be changed), the system raises an error if any code attempts to execute without a tenant context, and background jobs (report generation, calendar sync) carry and re-establish tenant context automatically.

## 4.2 Seat Enforcement

Each organization's subscription includes a defined number of staff and volunteer accounts. The authorization system enforces these limits — it is not possible to create accounts beyond the organization's allocation through any means, including administrative workarounds.

# 5 File Upload Security

StriveDB organizations upload sensitive documents including SANE reports, court orders, consent forms, and case photographs. File security is enforced through multiple layers:

Upload validation is mandatory on every model that accepts file attachments, verified at application startup. Validation includes a blocklist of dangerous file extensions (executables, scripts, archives), a 100 MB file size limit, and automated scanning of PDF files for embedded JavaScript — a technique used to deliver malware through seemingly benign documents. Malformed or encrypted PDFs that cannot be scanned are rejected.

Access control ensures uploaded files are accessible only to users who have authorization to view the parent record (e.g., a document attached to a survivor's case is accessible only to users who can access that survivor's record).

Storage security uses Amazon S3 with all public access blocked at the bucket level. Files are served via time-limited signed URLs that expire after 30 minutes — there are no permanent public links to any uploaded content.

# 6 Audit Trail

StriveDB maintains a comprehensive, tamper-evident audit trail across 56 data models. Every creation, modification, and deletion of a record is logged with the identity of the user who made the change, a timestamp, and a snapshot of the record's state.

Sensitive fields are excluded from the audit trail to prevent secondary exposure — passwords, encryption keys, and authentication tokens are never recorded in audit logs.

The audit trail supports:

- Investigating who accessed or modified a survivor's record and when
- Compliance with grant reporting requirements
- Forensic analysis in the event of a security incident
- Accountability for all data changes within the organization

# 7 Attack Prevention

## 7.1 Rate Limiting & Automated Attack Detection

StriveDB implements automated detection and blocking of common attack patterns at the application edge:

Behavioral blocklisting identifies and blocks IP addresses that send requests matching known attack signatures — path traversal attempts, admin panel probes, known vulnerability scanner patterns, and encoded attack payloads. Three suspicious requests within 10 minutes triggers a 10-minute IP ban.

Rate limiting enforces a ceiling of 200 requests per minute and 4,000 requests per hour per IP address, preventing both brute-force attacks and denial-of-service attempts.

## 7.2 Input Validation

All user input is validated and sanitized before processing:

- SQL injection prevention: All database queries use parameterized statements. Search inputs are sanitized using dedicated escaping functions. No raw SQL is constructed from user input.
- Cross-site scripting (XSS) prevention: All HTML output is automatically escaped. Rich text content is sanitized through a dedicated sanitization layer. No inline JavaScript is used in page templates.
- Cross-site request forgery (CSRF) protection: All state-changing requests require a cryptographic token verified by the server. Invalid tokens trigger a safe session cleanup and redirect to login.
- Mass assignment protection: The application raises an error if unexpected parameters are submitted, rather than silently ignoring them.
- Open redirect prevention: Redirect URLs are validated against an explicit allowlist of permitted paths, modules, controllers, and query parameters. Any redirect that does not match the allowlist is rejected.

## 7.3 Security Headers

The application sets security headers on all responses including Cache-Control directives that prevent browsers from caching sensitive pages, X-Frame-Options to prevent clickjacking by blocking the application from being embedded in other sites, HttpOnly and Secure flags on all session cookies, and HSTS at the infrastructure level to enforce HTTPS.

# 8 Privacy & Compliance

## 8.1 HIPAA Alignment

StriveDB is designed to support organizations that may be subject to HIPAA requirements. Our alignment with the HIPAA Security Rule includes:

**Administrative safeguards (45 CFR § 164.308):** Designated Security Officer, workforce security procedures with role-based access, MFA-enforced authentication, security awareness training program, documented incident response procedures, contingency planning with automated backups and Multi-AZ failover, and Business Associate Agreement with AWS.

**Technical safeguards (45 CFR § 164.312):** Unique user identification with mandatory MFA, audit controls via comprehensive change tracking across 56 models, person and entity authentication via multi-factor verification, and transmission security via TLS with HSTS enforcement.

**Physical safeguards (45 CFR § 164.310):** Managed by AWS for infrastructure (AWS maintains SOC 2, ISO 27001, and HITRUST certifications for their data centers). StriveDB maintains workstation security policies for the development team.

StriveDB is prepared to enter into Business Associate Agreements with customer organizations that require them.

## 8.2 VAWA, VOCA, and FVPSA Confidentiality Support

StriveDB's architecture supports the confidentiality requirements of the Violence Against Women Act, Victims of Crime Act, and Family Violence Prevention and Services Act through: complete organizational data isolation preventing any cross-organization data sharing, granular role-based access controls appropriate for separating volunteer access from clinical access from administrative access, individual record restriction capabilities for high-risk cases, safe-to-contact flags on survivor phone numbers and email addresses indicating whether contact at that number or address is safe (e.g., shared devices), and aggregate reporting capabilities (VOCA reports, demographic summaries) that satisfy grant requirements while minimizing individual-level data exposure.

## 8.3 Data Retention & Destruction

StriveDB maintains a formal Data Retention and Destruction Policy. Key provisions include: a 90-day grace period after subscription end for data export, documented destruction procedures with a Certificate of Destruction issued to departing organizations, soft-delete functionality allowing organizations to manage their own record lifecycle per their state laws and funder requirements, and compliance documentation retained for 6 years per HIPAA requirements.

## 8.4 Incident Response

StriveDB maintains a documented Incident Response Plan aligned with NIST SP 800-61 and HIPAA breach notification requirements. The plan includes incident classification and severity levels, containment and evidence preservation procedures, the HIPAA four-factor breach risk assessment, notification procedures with specific provisions for survivor safety (coordinating with affected organizations on safe notification methods that account for monitored mail and compromised email accounts), and post-incident review and continuous improvement processes.

## 8.5 Security Awareness Training

All StriveDB workforce members complete security awareness training covering HIPAA requirements, survivor data confidentiality, operational security practices, and incident recognition and reporting. Training is completed upon onboarding and refreshed annually, with quarterly security reminders and event-triggered retraining as needed. Training records are retained for 6 years.

### 8.6 Cyber Liability Insurance

StriveDB maintains active cyber liability insurance coverage.

## 9 Transparency: Areas of Ongoing Improvement

We believe in being straightforward about where we are strengthening our security posture:

**Content Security Policy:** We are implementing a comprehensive Content Security Policy header to provide an additional layer of defense against cross-site scripting attacks. The application's existing XSS protections (automatic output escaping, input sanitization, no inline scripts) provide strong protection today.

**Application-level encryption for narrative fields:** While all data is encrypted at rest via AWS RDS and in transit via TLS, we are evaluating application-layer encryption for narrative content fields (case notes, session notes, event descriptions) to provide an additional encryption boundary for the most sensitive free-text data.

**Automated data lifecycle management:** We are building automated retention enforcement to support organizations in managing record lifecycles per their state-specific requirements, including scheduled pruning of platform operational data (audit logs, analytics, session records) that has exceeded its retention period.

## 10 Questions We're Happy to Answer

We welcome security questions from prospective and current customers. Common topics include:

- **Penetration test results:** We can share a summary of findings and remediation status from our Accorian assessment upon request under NDA.
- **BAA execution:** We are prepared to enter into HIPAA Business Associate Agreements.
- **Specific technical controls:** We are happy to discuss any aspect of our security architecture in detail.
- **Data residency:** All data is stored in AWS US-West-2 (Oregon). We do not transfer data outside the United States.
- **Subprocessors:** Our infrastructure subprocessors include AWS (hosting, storage, compute, notifications, threat detection) and Sentry.

Contact us at info@strivedb.com or book a demo to discuss your organization's security requirements.