Data and Cyber Security

Incident AI
incident-ai.com

Incident AI is an AI-powered incident investigation assistant designed for high-risk industries such as mining, oil & gas, and construction. Ensuring the confidentiality, integrity, and availability of sensitive incident data is paramount. This white paper details the comprehensive security and safety measures integrated into Incident AI, covering architecture, data handling, encryption, access control, auditability, and compliance with global standards.

Introduction

As organizations increasingly leverage generative AI to streamline safety investigations, robust data security frameworks are essential to maintain trust and meet regulatory obligations. Incident AI combines advanced LLM-based analysis with stringent security controls to deliver actionable insights without compromising data privacy or system integrity.

System Architecture and Data Flow

Incident AI is hosted on Render's PaaS with microservices implemented in Node.js/Express and a React front end. Key architectural features include:

Frontend Security Stack

Web Application Firewall (WAF), DDoS protection, intrusion detection, and TLS 1.3 for all browser-to-server communications.

Backend Services

Express API service with input validation (Helmet, CORS, Morgan), JWT-based authentication with MFA (Auth0), and role-based access control.

Data Storage

PostgreSQL 16 on Render, with each customer assigned a separate schema; field-level encryption for sensitive fields. Automated daily backups and configurable disaster-recovery policies ensure data resilience.

External AI Integrations

Encrypted, zero-retention API calls to Anthropic Claude (reasoning), OpenAI Whisper (audio transcription), and Azure Computer Vision OCR (handwritten/text ingestion).

Encryption and Secure Transmission

In Transit

All API calls and external service interactions use HTTPS/TLS 1.3, ensuring end-to-end encryption.

At Rest

AES-256 encryption protects database files; field-level encryption secures personally identifiable information and incident details.

Key Management

Secrets and credentials are stored in Render's secure environment variables, with periodic rotation enforced by CI/CD pipelines.

Authentication, Authorization, and Identity Integration

Authentication

JWT tokens are issued upon login via Auth0, supporting SAML 2.0, OAuth 2.0, and OpenID Connect for federated SSO with Microsoft Entra ID or on-premises Active Directory.

MFA Enforcement

TOTP-based MFA is mandatory; policies can be delegated to the customer's IdP for centralized control.

Authorization

Role-Based Access Control (RBAC) enforced at both route and data layers. Express middleware validates JWT claims to scope data access by tenant and user role.

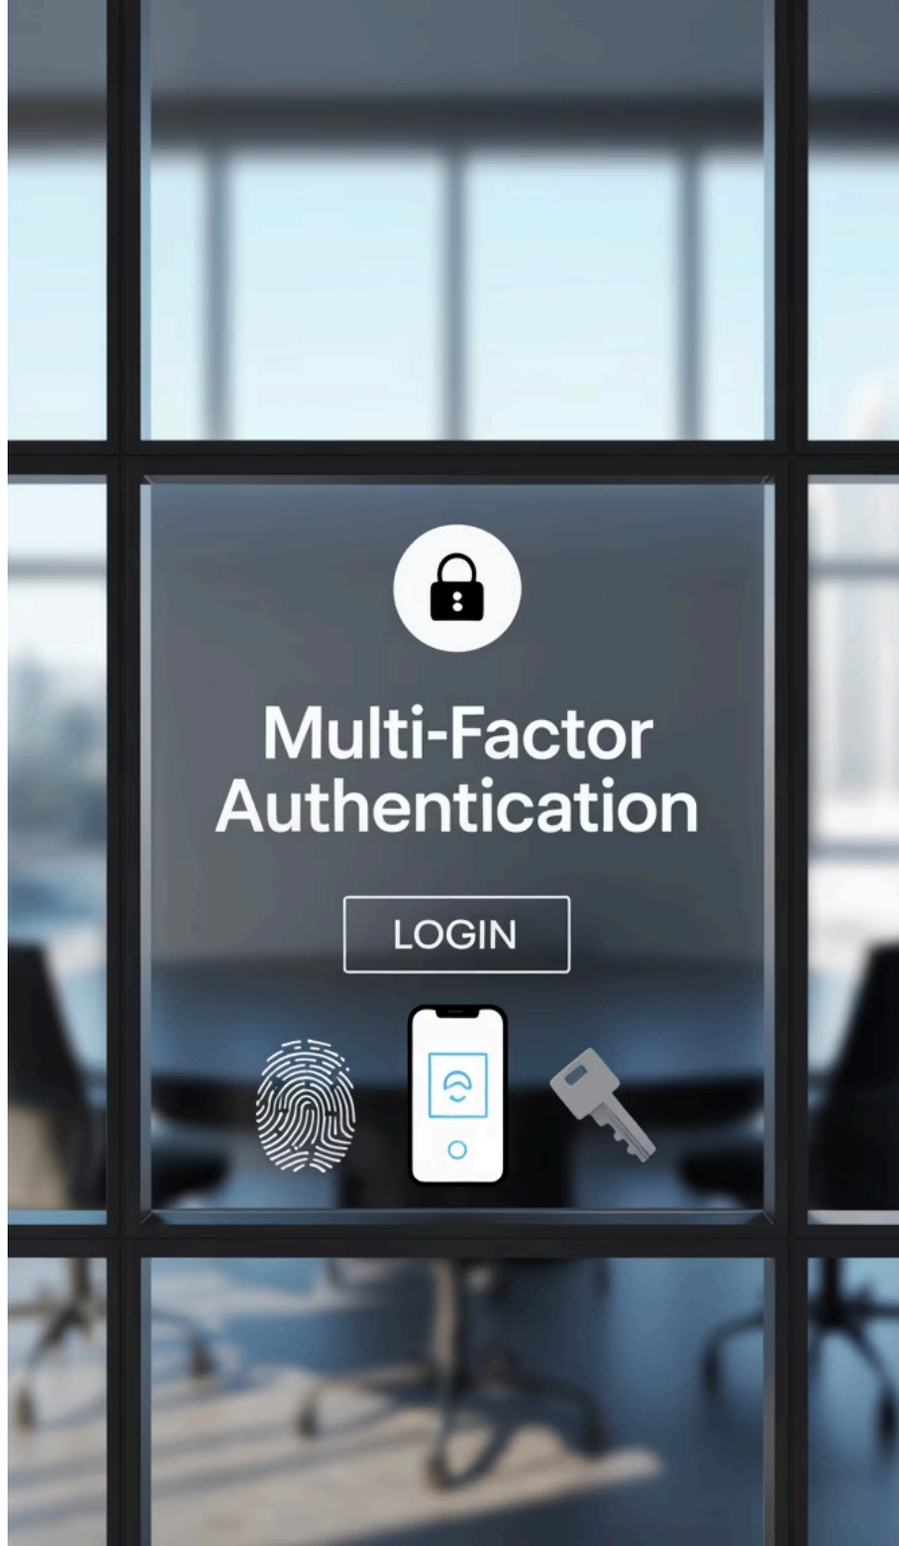

Data Processing, Retention, and Privacy

Zero-Retention Policy

Data processed by LLM components resides only in volatile memory; no transcripts or input data are stored for model training or reuse.

Third-Party Compliance

External AI services comply with GDPR, SOC 2 Type 2/Type 1 (Anthropic), SOC 3, CCPA, HIPAA, and ISO/IEC 27018 standards.

Subscription Data

Customer-uploaded incident files and AI-generated outputs are stored only within the customer's schema and deleted upon case closure or subscription termination per automatic deletion policies.

Audit Logging and Monitoring

Structured Logging

Winston captures all authentication events, data edits, and system errors with user- and tenant-level metadata. Morgan logs HTTP request metrics for performance monitoring.

Audit Trails

All access and modification actions (e.g., prompt updates, case changes) are logged to support post-incident review and compliance audits.

Health Monitoring

Render's built-in metrics (CPU/memory usage, request latency) and planned Prometheus integration ensure proactive incident detection.

Backup, Disaster Recovery, and Availability

Automated Backups

Daily PostgreSQL backups with seven-day retention window; manual on-demand backups supported.

Uptime Commitments

98% service availability guaranteed, with service credits for underperformance.

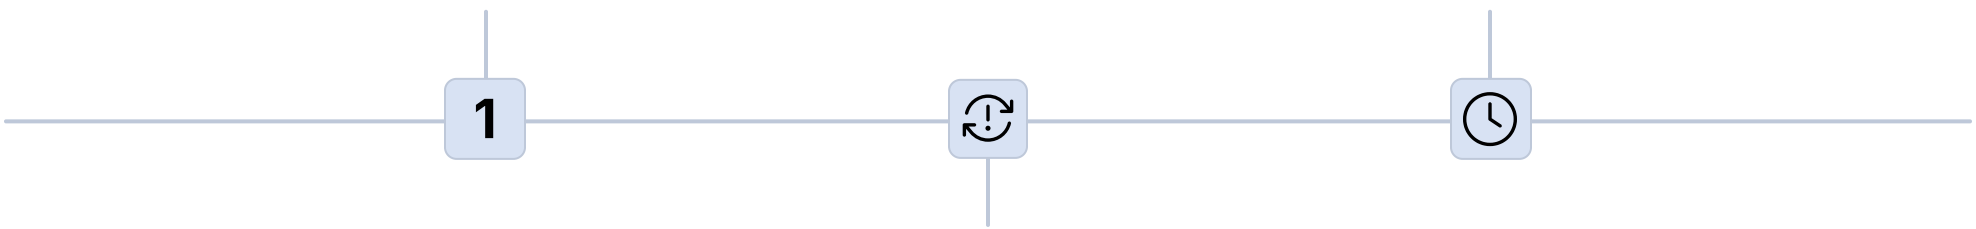

Disaster Recovery

Multi-region failover options via Render's infrastructure; recovery time objectives (RTO) and recovery point objectives (RPO) defined per customer SLAs.

Compliance and Governance

Regulatory Alignment

Incident AI adheres to GDPR, CCPA, HIPAA, POPIA, and Australian Privacy Principles. Data processing agreements and breach-notification procedures comply with applicable laws.

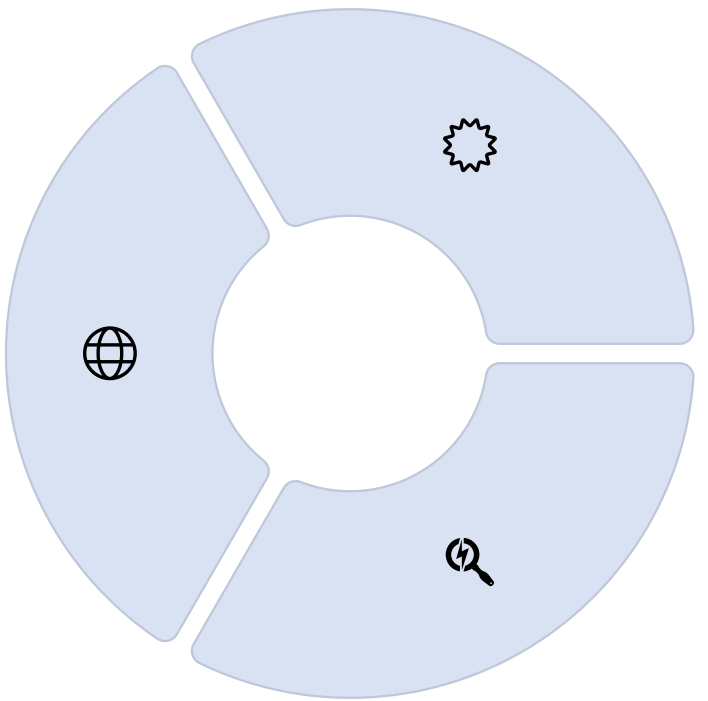

Security Standards

SOC 2 Type 2 certification of platform components, ISO/IEC 27001 readiness, and CSA Star compliance for cloud services.

Third-Party Audits

Optional on-premises hosting undergoes independent security assessments; Render's regular penetration testing and compliance audits bolster overall security posture.

Conclusion

Incident AI delivers powerful AI-driven incident investigation capabilities underpinned by a rigorous security and safety framework. By combining modern encryption standards, federated identity management, zero-retention AI processing, and comprehensive auditability, Mine Guard AI ensures that organizations can harness generative AI with confidence and compliance in high-risk environments.

MINE GUARD AI