GUARDIAN PLATFORM

Security Policy

SaaS Compliance Infrastructure

VERSION	EFFECTIVE DATE	CLASSIFICATION
2.0 — Public Release	March 2026	PUBLIC

> *"Bridging the gap between Algorithmic Performance and EU Compliance."*

Nordic AI Integrity ApS · CVR 46224582 · Copenhagen, Denmark
nordicaiintegrity.dk · thomas@nordicaiintegrity.dk

Table of Contents

Guardian Security Policy v2.0 — Public Document

SECTION 01

Introduction & Scope

Nordic AI Integrity ApS operates Guardian, a SaaS platform providing EU AI Act compliance monitoring for high-risk AI systems under Annex III. Guardian processes compliance metrics — drift scores, fairness indices, and model metadata — to help regulated organisations maintain continuous compliance.

> **Core Principle:** *Guardian is designed to operate on aggregated compliance metrics rather than raw personal data. Clients should avoid submitting raw personal data through the platform or API. This design choice fundamentally reduces data protection risk.*

This policy covers:

- The Guardian SaaS platform and all associated APIs
- Internal systems, tools, and administrative access
- Client data (compliance metrics, model metadata, audit reports)
- Infrastructure hosting and third-party service providers
- Employee and contractor security obligations

Security objectives:

- **Confidentiality** — Client compliance data is protected against unauthorized access
- **Integrity** — Compliance metrics and audit trails remain accurate and tamper-proof
- **Availability** — Guardian is designed for high availability and continuous monitoring
- **Compliance** — Alignment with GDPR, EU AI Act, and industry best practices

SECTION 02

Zero Trust Architecture

Guardian is built on a Zero Trust security model guided by the principle "never trust, always verify." Every access request is authenticated, authorized, and validated regardless of network origin.

2.1 Core Zero Trust Principles

- **Verify Explicitly** — Every API call and user session is authenticated with a valid bearer token. No implicit trust is granted based on network position or prior authentication.
- **Least Privilege Access** — Users and services receive the minimum permissions required for their function. Permissions are reviewed quarterly.
- **Assume Breach** — Systems are designed under the assumption that any component may be compromised. Blast radius is minimised through micro-segmentation and strict tenant isolation.

2.2 Implementation

- Strict tenant isolation: each client organisation's data is logically separated at the database level using organisation-scoped queries. Cross-tenant access is prevented by organisation-scoped access controls, application-layer authorization, and strict tenant isolation patterns.
- API authentication via secure bearer tokens with automatic expiration and rotation policies.
- Session management with NextAuth.js: short-lived sessions, secure cookies, CSRF protection.
- Network-level controls: all internal service communication is encrypted and authenticated.
- Continuous monitoring of access patterns with automated anomaly detection.

SECTION 03

Multi-Factor Authentication (MFA)

Multi-Factor Authentication is mandatory for all access to Guardian systems, both for internal team members and for client administrators with elevated privileges.

3.1 MFA Requirements

User Type	MFA Requirement	Methods
Internal Team (all)	Mandatory	TOTP / Hardware key
Guardian Admin	Mandatory + review	TOTP + hardware key
Client Owner	Mandatory	TOTP / Authenticator app
Client Viewer	Recommended	TOTP / Authenticator app
API Access	Token + IP whitelist	Bearer token + allowlist

3.2 Password Policy

- Minimum 12 characters with complexity requirements (upper, lower, number, special)
- Passwords are hashed using bcrypt with a minimum cost factor of 12
- Credential stuffing protection: rate limiting and account lockout after 5 failed attempts
- Mandatory password rotation every 90 days for internal admin accounts

3.3 Single Sign-On (SSO)

Guardian is designed to integrate with enterprise identity providers (IdPs) through industry standards such as SAML 2.0 and OpenID Connect. Support for SSO and Just-In-Time (JIT) user provisioning is on the product roadmap to allow clients to enforce their own identity, MFA, and lifecycle policies centrally.

Status: Planned — Product roadmap

SECTION 04

Role-Based Access Control (RBAC)

Guardian enforces strict Role-Based Access Control to ensure that every user has access only to the resources required for their function. Roles are defined at the platform level and scoped to individual client organisations.

4.1 Role Definitions

Role	Scope	Capabilities
GUARDIAN_ADMIN	Platform-wide	Full platform administration, manage all client organisations, import data, manage users
CLIENT_OWNER	Organisation	View and manage own organisation's AI systems, update alert statuses, upload documents, manage team
CLIENT_VIEWER	Organisation (read)	Read-only access to own organisation's dashboards, reports, and compliance status

4.2 Access Governance

- Quarterly access reviews: all user permissions are audited by the security owner
- Immediate revocation upon role change or departure
- Separation of duties: no single individual can both approve and execute critical changes
- Complete audit trail: every permission change is logged with timestamp and actor

SECTION 05

SaaS Security Posture Management (SSPM)

Nordic AI Integrity maintains a comprehensive SaaS Security Posture Management programme to monitor, assess, and harden the security configuration of all SaaS tools and services used in our operations and platform delivery.

5.1 SaaS Inventory & Configuration

- Centralised inventory of all SaaS services (GitHub, Vercel, Neon DB, monitoring tools)
- Baseline security configuration for each service, reviewed quarterly
- Automated alerts for configuration drift or unauthorised changes
- Shadow IT detection: any new SaaS adoption must be approved by the security owner

5.2 Continuous Monitoring

- Centralised logging of all authentication events and administrative actions
- Real-time alerting on suspicious activity (unusual login locations, privilege escalation)
- Automated vulnerability scanning of dependencies and container images
- Regular penetration testing of external-facing services

5.3 Third-Party Risk Management

- All third-party SaaS providers are assessed for SOC 2 or equivalent certifications
- Data processing agreements (DPAs) in place with all sub-processors
- Annual review of sub-processor security posture and compliance status
- Contractual right to audit and breach notification obligations

Current sub-processors include: Vercel (hosting), Neon (database), GitHub (source control). A complete list is available at nordicaiintegrity.dk/subprocessors.

SECTION 06

Data Encryption

All data handled by Guardian is encrypted both in transit and at rest, using industry-standard cryptographic protocols. Our encryption strategy ensures that even in a breach scenario, data remains unreadable.

6.1 Encryption Standards

Layer	Standard	Details
Data in Transit	TLS 1.3	All API calls, web traffic, and internal communications. HSTS enforced. TLS 1.0/1.1 disabled.
Data at Rest	AES-256	Database encryption via cloud provider managed keys. Backup encryption enabled.
API Tokens	SHA-256 + salt	Tokens are hashed before storage. Raw tokens are never persisted.
Passwords	bcrypt (cost 12)	Adaptive hashing function. Plaintext passwords are never stored or logged.

6.2 Key Management

- Encryption keys are managed by the cloud provider's key management service (KMS)
- Key rotation occurs automatically on a scheduled basis
- Access to key management is restricted to designated security personnel only
- Key usage is fully audited and logged

> **Zero Raw Data Architecture:** *Guardian processes compliance metrics only — drift scores, fairness indices, sample sizes. We never receive, store, or process raw personal data from client AI models. This fundamentally limits our data exposure surface.*

SECTION 07

Incident Response Plan

Nordic AI Integrity maintains a formal incident response plan to ensure rapid, coordinated, and transparent handling of any security event affecting Guardian or client data.

7.1 Severity Classification

Level	Definition	Response	Notification
CRITICAL	Active data breach, material compromise, or full service outage	Within 1 hour	Immediate internal escalation and client notification as applicable
HIGH	Significant vulnerability or partial service degradation	Within 4 hours	Client notification within 24 hours if impacted
MEDIUM	Suspected anomaly, minor vulnerability, configuration issue	Within 24 hours	Internal escalation; client if impacted
LOW	Informational finding, policy deviation, non-urgent improvement	Within 72 hours	Logged internally

7.2 Response Phases

1. **Detection & Triage** — Automated monitoring systems and manual reports feed into our incident triage process. Every alert is classified within the SLA defined above.
2. **Containment** — Affected systems are isolated immediately. Access tokens are revoked. Forensic evidence is preserved before any remediation.
3. **Eradication & Recovery** — Root cause is identified and eliminated. Systems are restored from verified backups. Integrity checks confirm clean state.
4. **Communication** — Affected clients are notified per GDPR Article 33/34 timelines. Datatilsynet (Danish DPA) is notified within 72 hours if personal data is affected.
5. **Post-Incident Review** — A formal post-mortem is conducted within 5 business days. Lessons learned are documented and corrective actions are implemented.

SECTION 08

Compliance & Regulatory Alignment

Guardian is designed to operate within a complex regulatory environment. Our compliance programme is continuously evolving to meet current and upcoming requirements.

8.1 Regulatory Frameworks

Framework / Control	Status	Notes
GDPR	Current	Metrics-only approach, EU hosting, DPAs with all sub-processors
EU AI Act alignment	Current	Supports Art. 9, 11, Annex IV workflows
SSO (SAML / OIDC)	Planned	Product roadmap
SOC 2 Type II	Planned	Target 2027
ISO 27001	Planned	Target 2027, post-SOC 2

8.2 Data Residency

All Guardian data is processed and stored within the European Union. Our primary database is hosted on Neon (PostgreSQL) in the EU-Central region. Application hosting is on Vercel with EU edge deployment. No client data is transferred outside the EEA.

8.3 Data Retention & Deletion

Compliance metrics, audit logs, and related Guardian data are retained for a standard period of 24 months unless otherwise agreed in a data processing agreement. Backups follow the same retention window and are automatically purged after expiry.

Upon contract termination or written request from a client, Nordic AI Integrity will delete or irreversibly anonymise all Guardian data for that client within 30 days, subject to any mandatory legal retention requirements.

Note: Where AI explanation features are enabled, prompt and output retention periods may be shorter than the standard platform retention window and should be defined in product-specific settings or contractual

documentation.

SECTION 09

Infrastructure & Hosting

Guardian is deployed on modern, cloud-native infrastructure designed for resilience, scalability, and security.

9.1 Architecture Overview

Component	Provider	Security Measures
Application	Vercel (Next.js 14)	Edge deployment, automatic HTTPS, DDoS protection, serverless architecture
Database	Neon (PostgreSQL 16)	EU-hosted, encrypted at rest, connection pooling, SSL-required connections
Authentication	NextAuth.js	Credential provider with bcrypt hashing, secure sessions, CSRF protection
ORM / Data Layer	Prisma	Parameterised queries (SQL injection prevention), strict type safety
Source Code	GitHub (private)	Branch protection, required reviews, signed commits, dependency scanning

9.2 Application Security

- Secure coding practices with TypeScript strict mode and ESLint security rules
- Parameterised queries via Prisma ORM — eliminating SQL injection vectors
- Input validation on all API endpoints with server-side type checking
- HTTP security headers: HSTS, X-Content-Type-Options, X-Frame-Options
- Automated dependency vulnerability scanning via GitHub Dependabot
- Regular code reviews with security-focused checklists

SECTION 10

Policy Governance & Review

This security policy is a living document owned by the CEO and reviewed regularly to reflect evolving threats, regulatory changes, and business growth.

This document is reviewed at least annually and may be updated earlier to reflect infrastructure, product, or regulatory changes.

10.1 Review Schedule

- **Annual comprehensive review** — Full policy revision including controls assessment, threat landscape update, and regulatory alignment check
- **Quarterly access review** — All user permissions, API keys, and third-party access are reviewed and revalidated
- **Event-triggered review** — Policy is reviewed immediately after: a security incident, significant infrastructure change, new regulatory requirement, or major client onboarding

10.2 Policy Ownership

Responsibility	Owner
Overall security policy	Thomas Noba, CEO
Technical implementation	Dr. OJ Akintande, Technical Lead
Compliance & regulatory	Joris Cappa, COO
Incident response lead	CEO until a dedicated security lead is appointed

> **Contact:** *For security inquiries, vulnerability reports, or to request a detailed security questionnaire response, contact thomas@nordicaiintegrity.dk or joris@nordicaiintegrity.dk*

Security Contact: security@nordicaiintegrity.dk
Vulnerability Disclosure: Responsible disclosure accepted
Security Summary:nordicaiintegrity.dk/security
Subprocessors:nordicaiintegrity.dk/subprocessors

Thomas Noba

Co-founder & CEO, Nordic AI Integrity ApS

Effective: March 2026