Security & Compliance Framework

Elis AI Technical Documentation

Generated: 2026-05-19 12:44 UTC

Docs: CIS Hardening Guide

CIS Hardening Guide

This guide documents baseline hardening controls for production deployment and the CI scan workflow for drift visibility.

Current Runtime Hardening

The production stack enforces key container hardening controls in docker-compose.prod.yml:

- read-only filesystems on sensitive services.
- no-new-privileges security option.
- Linux capabilities dropped (`cap_drop: ALL`) on hardened services.
- tmpfs mounts for transient writable paths.
- segmented Docker networks separating frontend/backend/data/miners.

CI Security Checks

The platform quality workflow runs in .github/workflows/ci-platform-gate.yml and includes:

- Bandit static security scan for backend code.
- Python dependency vulnerability scan via pip-audit.
- Frontend dependency vulnerability scan via npm audit.
- OpenSCAP baseline scan with generated ARF + HTML artifacts.

OpenSCAP artifacts:

- `openscap-results.arf`
- `openscap-report.html`

These are uploaded in CI for security review and remediation tracking.

Baseline Drift Detection

Use Step 10 change tracking endpoints as baseline anchors:

- `GET /api/system-admin/change-log`
- `GET /api/system-admin/baseline`

Review drift by comparing:

- deployed image digest and environment snapshot from `/baseline`
- expected compose hardening settings from source control
- OpenSCAP report findings from CI artifacts

Recommended Review Cadence

- On every merge to main/master: review security scan job output.
- Weekly: compare latest OpenSCAP report against previous run.
- Monthly: verify runtime compose hardening settings match documented baseline.

Docs: Security Training Program

Security Awareness Training Program

Purpose

This document defines the security awareness and training program for all personnel with access to the Elis AI platform, covering HIPAA, ITAR, NERC CIP, FISMA, and PCI-DSS training requirements.

Training Modules

Module 1: General Security Awareness (All Personnel)

- Information security policy overview
- Password management and authentication best practices
- Social engineering and phishing awareness
- Incident reporting procedures
- Acceptable use of technology and data
- Physical security awareness (for on-premise deployments)

Module 2: HIPAA / PHI Handling (Healthcare Deployments)

- Definition of Protected Health Information (PHI)
- Minimum Necessary Standard
- PHI access, use, and disclosure requirements
- Breach reporting obligations and HITECH notification rules
- Sanctions for HIPAA violations

Module 3: ITAR / Export Control (Defense Deployments)

- ITAR/EAR scope, definitions, and US Munitions List overview
- Penalties for export control violations
- Proper handling and marking of controlled technical data
- US-person verification and access restrictions
- Technology Control Plan (TCP) awareness

Module 4: NERC CIP (Energy Sector Deployments)

- BES Cyber System awareness
- Physical and electronic security perimeter policies
- Proper handling of BES Cyber System Information (BCSI)
- Incident reporting for cyber security events
- CIP policy compliance requirements

Module 5: FISMA / Federal (Government Deployments)

- Federal information security requirements
- Insider threat awareness
- Incident reporting to US-CERT/CISA
- Handling of Controlled Unclassified Information (CUI)
- Security responsibilities by role

Module 6: PCI-DSS (Payment Processing Environments)

- Cardholder data handling and protection
- PCI scope and SAQ-A eligibility
- Acceptable use policies for systems in scope
- Incident response for payment data breaches

Training Schedule

Audience	Frequency	Modules
All personnel	Upon hire + annually	Module 1
PHI-access personnel	Upon assignment + annually	Modules 1, 2
ITAR-access personnel	Upon assignment + annually	Modules 1, 3
BES Cyber System personnel	Quarterly awareness + annually	Modules 1, 4
Federal system personnel	Upon assignment + annually	Modules 1, 5
Payment environment personnel	Upon hire + annually	Modules 1, 6
Privileged users (admin+)	Upon assignment + annually	Module 1 + role-specific

Training Delivery

- Online self-paced modules with completion tracking
- Quarterly phishing simulation exercises
- Annual tabletop incident response exercises
- Role-based hands-on security training for system administrators

Record Keeping

- Training completion records retained per regulation:
- HIPAA: 6 years
- ITAR: 5 years
- NERC CIP: 3 calendar years
- FISMA: 3 years
- PCI-DSS: 1 year (recommended 3 years)
- Records include: personnel name, training date, module(s) completed, assessment score

Compliance Tracking

- Training completion tracked via company_audit_log with `training_completed` event type
- Overdue training alerts generated via compliance alerts API
- Access restricted for personnel with overdue mandatory training (enforced via clearance_level)

Docs: SQL Privacy Verification

SQL privacy and UX verification (manual)

Use this checklist after changing SQL tooling, traces, or chat blocks.

Automated

```
python -m pytest backend/tests/test_sql_query_tool.py backend/tests/test_sql_tool_loop_privacy.py
backend/tests/test_sql_blocks_merge.py -q
```

Optional broader run:

```
python -m pytest backend/tests/test_orchestration.py -k sql --tb=short -q
```

Manual (auth → message)

1. Sign in and open a conversation with a linked company database (or dev SQL fixtures).
2. Ask a question that triggers `sql_query` (e.g. row counts or list tables).
3. Confirm the assistant message shows **Summary** and **Data** tabs when tabular results return.
4. On **Data**, confirm column sort, filters, and **Visualize** / **Ask AI to chart** behave as expected.
5. Open trace / observability (if available) and confirm SQL trace rows contain no raw connection strings or passwords.

Environment flags to spot-check

- `SQL_SCHEMA_SNAPSHOT_STRICT=1` — mismatched `schema_snapshot_id` should error with `schema_snapshot_mismatch` / `SCHEMA_SNAPSHOT_MISMATCH`.
- `TOOL_PRIVACY_BROKER_ENABLED=1` — SQL answer text may be redacted when broker is active.

Docs: Incident Response Plan

Incident Response Plan

This document defines the operational workflow for breach and cyber incident response.

Scope

Applies to security incidents affecting confidentiality, integrity, availability, or tenant isolation.

Roles

- Incident Commander (IC): owns severity assignment, escalation, and closure approval.
- Technical Lead (TL): drives containment, forensics, and remediation.
- Communications Lead (CL): owns internal and external notifications.
- Compliance Lead (optional): maps incident actions to HIPAA/NERC reporting obligations.

Severity Levels

- critical: active or likely compromise, major data exposure, cross-tenant impact.
- high: probable compromise with sensitive data risk.
- medium: suspicious activity requiring investigation.
- low: low-confidence anomaly with no immediate impact.

Detection Sources

- System-admin anomaly endpoint: GET /api/system-admin/breach-incidents/anomalies?company_id=
- SIEM stream (Fluent Bit -> Azure Log Analytics)
- Manual analyst reports

Workflow

1. Detect
   - Gather anomaly evidence and affected users.
   - Create incident via POST /api/system-admin/breach-incidents.
2. Assess
   - Set severity and status to assessing.
   - Estimate affected records count.

3. Contain
   - Revoke tokens and credentials as needed.
   - Restrict access paths and isolate impacted systems.
   - Update incident status to contained.
4. Notify
   - HIPAA: notification deadline is detected_at + 60 days.
   - NERC CIP reportable incidents: initiate 1-hour notification path.
   - Set status to notified and record notified_at.
5. Recover and Close
   - Validate controls are restored.
   - Document root cause and corrective actions.
   - Set status to resolved after IC sign-off.

API Runbook

- Create incident:
- POST /api/system-admin/breach-incidents
- List incidents:
- GET /api/system-admin/breach-incidents
- Update incident:
- PATCH /api/system-admin/breach-incidents/
- Scan anomalies:
- GET /api/system-admin/breach-incidents/anomalies?company_id=

Audit and Evidence

For each incident, preserve:

- timeline of status changes
- affected systems and records estimate
- notification timestamps and recipients
- remediation tasks and verification results

Escalation Matrix

- critical: page IC and TL immediately.

- high: page TL, notify IC within 15 minutes.
- medium: assign on-call security engineer within 1 hour.
- low: queue for business-hours triage.

Docs: BAA Registry

BAA Registry

Purpose

Track Business Associate Agreement (BAA) coverage for regulated deployments.

Registry

Provider	Service Scope	BAA Status	Notes
Microsoft Azure	Hosting, managed database, networking, monitoring	In progress / required	Enterprise agreement path for BAA execution and documentation.
Stripe	Billing and subscription events	Required	Confirm BAA or equivalent healthcare data handling terms if PHI-related flows include billing metadata.
Email Provider (SMTP/ACS)	Authentication and notification delivery	Required	Confirm BAA-capable plan and retention policy alignment.
AI Model Provider (BYOK mode)	Inference APIs	Customer-managed	In BYOK deployments, customer owns direct provider agreement and BAA obligations.

Data Flow Notes

- Dedicated/BYOK deployments reduce shared sub-processor exposure.
- Maintain this document as part of procurement and compliance review cycles.

Review Cadence

- Quarterly: verify contract status and expiry dates.
- On provider change: update registry before production use.

Docs: FISMA Security Categorization

FISMA Security Categorization (FIPS 199)

System Context

Elis AI is a multi-tenant orchestration platform handling conversational inputs, uploaded organizational documents, and generated model outputs.

Information Type Categorization

Information Type	Confidentiality	Integrity	Availability	Rationale
User prompts and conversation metadata	Moderate	Moderate	Moderate	May contain sensitive operational context and potentially regulated data.
Company uploaded documents and extracted knowledge	Moderate	Moderate	Moderate	Organizational proprietary content; integrity impacts decision quality.
AI-generated responses	Low	Moderate	Moderate	Typically derived output, but integrity and availability impact operations.
Authentication/session/security logs	Moderate	Moderate	Moderate	Required for incident response and audit accountability.

Resulting Security Category

Overall provisional categorization: **Moderate** baseline.

- Confidentiality: Moderate
- Integrity: Moderate
- Availability: Moderate

Control Implications

A Moderate baseline requires layered safeguards including:

- strong access control and least privilege enforcement
- audit logging and monitoring
- encryption for sensitive data at rest and in transit
- incident response and contingency planning

Docs: FISMA SSP

System Security Plan (NIST SP 800-18)

1. System Identification

- System name: Elis AI Platform
- System owner: Platform Operations / Security Team
- Authorization boundary: frontend, backend API/orchestrator, database, cache, miners, supporting monitoring and CI controls

2. System Environment

- Deployment model: shared or dedicated organizations
- Primary controls: tenant scoping, RBAC, clearance controls (when enabled), encryption controls, network segmentation, CI security scans

3. Control Families (Implementation Snapshot)

AC - Access Control

- authenticated access required for protected APIs
- role-based enforcement for organization operations
- optional clearance and compartment controls for classified data paths

AU - Audit and Accountability

- `company_audit_log` records organization-level actions
- `system_change_log` tracks baseline and change events
- system-admin notification and compliance alert surfaces for operational visibility

CM - Configuration Management

- baseline snapshot endpoint (`/api/system-admin/baseline`)
- change log endpoint (`/api/system-admin/change-log`)
- CI security quality gate for code and dependency scanning

IR - Incident Response

- breach incident workflow persisted in `breach_incidents`

- anomaly detection service for suspicious audit patterns
- incident runbook in docs/incident-response-plan.md

SC - System and Communications Protection

- segmented production networks in compose deployment
- optional outbound egress lock (`NO_EXTERNAL_EGRESS`)
- geo-fence middleware support (`GEO_FENCE_REGIONS`)

CP - Contingency Planning

- dedicated deployment lifecycle scripts
- disaster recovery guidance in docs/disaster-recovery.md
- rollback procedure in scripts/rollback.sh

4. Ongoing Assessment

- periodic compliance checks and alerts via system-admin compliance operations
- generated compliance reports for framework-level evidence snapshots

5. Attachments

- Security Categorization: docs/fisma/security-categorization.md
- POA&M: docs/fisma/poam.md

Docs: FISMA POAM

Plan of Action and Milestones (POA&M)

Open Items

ID	Gap	Severity	Planned Action	Owner	Target
POAM-001	Full continuous compliance scheduler (automated periodic check execution)	Medium	Add periodic scheduler to run compliance checks for all companies every 6 hours and persist execution telemetry	Backend Platform	2026-05-15
POAM-002	End-to-end geo-fence validation in staging with representative IP test harness	Medium	Add staging validation suite for non-US and US source simulation	Security Engineering	2026-05-01
POAM-003	Expanded OpenSCAP profile tailoring for container-specific baseline	Low	Add policy-tailored profile and remediation tracking docs	DevSecOps	2026-05-20

Completed Items

ID	Item	Completion Date
C-001	Breach incidents table + APIs + anomaly detection service	2026-04-11
C-002	Compliance operations APIs for alerts/reports	2026-04-11
C-003	Step 10 change tracking endpoints and baseline snapshot	2026-04-11
C-004	File integrity monitoring sidecar integration	2026-04-11

Site: /docs

Documentation

Learn how the Elis AI distributed intelligence network processes your requests.

How Elis AI Works

Elis AI is a distributed intelligence network that breaks complex questions into smaller, focused tasks and routes them to specialized AI agents running across a decentralized pool of compute miners. Instead of relying on a single large model, every request passes through a dynamic pipeline that picks the best-fit agent, dispatches inference to the best available miner, and returns a fully traced answer — so you always know how a response was produced.

🤖 Elis AI Mode

Automatically builds and reuses specialized agents. Each agent runs a multi-step workflow — expertise lookup, context retrieval, execution — using the best available model tier for each step.

🔓 Elis Unlocked

Advanced multi-pass workflow with planning, fact-checking, and iterative refinement. Best for complex tasks that need verification and deep research.

⛏️ Miners

Miners host local LLMs and earn tokens for verified responses. The network supports multiple model tiers for speed and quality tradeoffs.

🛡️ Verification

Every inference step is traced end-to-end. The scheduler selects the best available miner using weighted round-robin, trust tiers, and load awareness — ensuring fair distribution and reliable execution.

Why the network stays sustainable

Elis is built on a token economy where compute contributors are rewarded for useful work. Miners that return high-quality results quickly earn more assignments. Because miners compete regionally, users get served by nearby high-performing nodes — improving latency and experience.

For Consumers

Contribute background compute to earn tokens that offset usage costs.

Background contribution earns tokens.

Tokens are spent on model and tool usage.

Quality verification keeps rewards fair.

For Enterprise

End-to-end encryption, automated PII redaction, audit trails, and self-hosted deployment options. See Enterprise section →

Process Flow

1. User submits prompt → 2. Router chain selects agent → 3. Agent builds workflow → 4. Miners run inference → 5. Quality verification → 6. Traced response returned

Ready to try it?

Create an account and start asking questions — every answer comes with a full trace.

Get Started → View Whitepapers