FROM COMPLIANCE TO CONTROL: GOVERNING GENAI THROUGH THE CISO FUNCTION

Ilona Veitaitė[0000-0001-9046-0788], Rasa Brūzgienė[0000-0002-0816-8700], Ieva Andrijauskaitė[0009-0007-6315-0427] and Donatas Misiūnas[0000-0002-8253-684X]
Institute of Social Sciences and Applied Informatics, Vilnius University
ilona.veitaite@knf.vu.lt

DATA ANALYSIS METHODS FOR SOFTWARE SYSTEMS

VILNIAUS UNIVERSITETAS · 1579 · UNIVERSITAS VILNENSIS

Kaunas Faculty

Why GenAI needs its own governance

GENAI BEHAVES DIFFERENTLY

- **Plausible but wrong**
- **Open to manipulation**
- **Drifts over time**

FOUR EU LAWS, ONE AI PROCESS

AI Act
Risk-based AI rules

GDPR
Personal data protection

CRA
Secure digital products

NIS2
Cyber risk and incidents

The approach: the Matrix-Plus-Map

THE MATRIX
What must the CISO do?

12 CISO functions linked to EU law and ISO standards

THE MAP
How much can AI do?

Each function placed on the human-AI boundary

WHAT WAS DONE

1 Mapped functions → **2** Classified AI share → **3** Designed process → **4** Tested two cases

How the CISO-led workflow runs

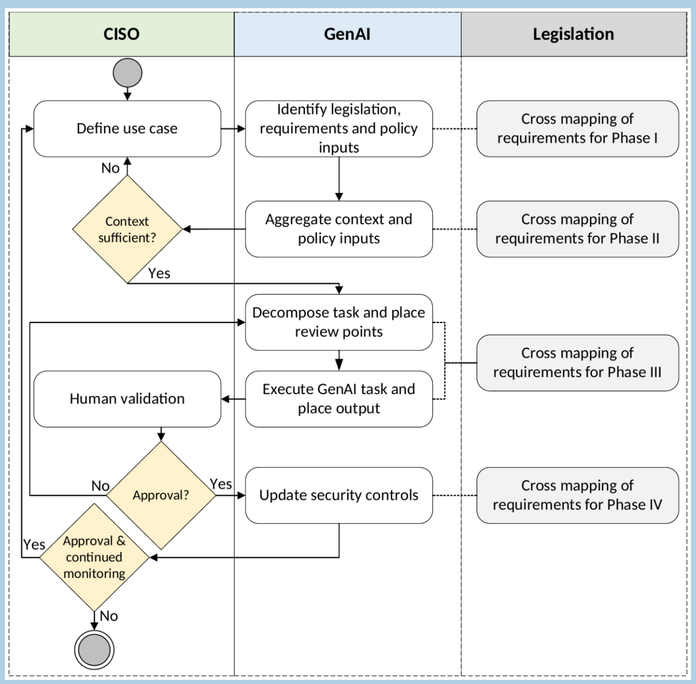

1. **CISO** defines the use case
2. **GenAI** finds the relevant law and gathers context
3. **GenAI** splits the task, sets review points and does the work
4. **Human** validates: only approved output updates the security controls

Legislation: cross-mapped in every phase (I–IV)

The Matrix: 12 CISO functions

Threats and risk	Compliance and maturity	Incidents and recovery	Emerging threats and validation	Access and readiness
F1 Threat analysis	**F2** Policy mapping	**F6** Incident response	**F8** LLM threat detection	**F11** Adaptive access control
F5 Alert correlation	**F3** Maturity assessment	**F7** Recovery guidance	**F9** Testing and validation	**F12** Crisis training scenarios
F10 Vulnerability priorities				
F4 Executive briefings				

Applying M&M in five steps

1. **Regulatory mapping** — Laws and standards per function
2. **Control mapping** — ISO/IEC 27002 controls, organizational or technical
3. **AI augmentation analysis** — Automate, augment, or keep human?
4. **Human oversight boundary** — Decision points and escalation paths
5. **Validation** — Test in a real organization

In daily operation, security controls change **only after human approval**.

One function, five frameworks

F6 · Incident response
a single CISO function

- **NIS2** — Incident handling, business continuity
- **AI Act** — Risk management, serious-incident reporting
- **CRA** — Reporting exploited vulnerabilities and incidents
- **ISO 37000** — Accountability, risk governance
- **ISO/IEC 27001** — Incident response, logging and monitoring controls

Every function traces to several instruments at once: **compound accountability**

Use case A: University

GenAI classifies research data sensitivity and drafts data management plans

REGULATION
NIS2 · GDPR · ISO 27001 · AI Act: limited risk

CISO FUNCTIONS
F1 · F2 · F9 · F11

AI AUGMENTATION
~72%
average over four AI-supported tasks

4 HUMAN DECISION POINTS
1. Reclassify data to “restricted”
2. Accept the classification error rate
3. Approve model updates before deployment
4. Grant privileged access exceptions

The Map: three questions per function

1 REVERSIBILITY
Can it be undone?
- AI — Reversible, low-impact tasks
- Human — Irreversible legal, financial or reputational impact

2 ACCOUNTABILITY
Who is accountable?
- AI — Predefined compliance checks
- Human — Personal liability and human oversight duties

3 CONTEXTUAL JUDGMENT
Does it need judgment?
- AI — Pattern recognition at scale
- Human — Ethics, trade-offs, interpreting the law

Use case B: FinTech

GenAI detects fraud in real time and scores customer risk

REGULATION
NIS2 · DORA · CRA · ISO 27001 · AI Act: high risk

CISO FUNCTIONS
F1 · F5 · F6 · F8 · F10 · F11

AI AUGMENTATION
~74%
average over five AI-supported tasks

5 HUMAN DECISION POINTS
1. Block a transaction above €50,000
2. Model drift beyond ±5%: continue or retrain
3. Notify the CSIRT and financial authority
4. Override an AI “legitimate” label
5. Emergency break-glass model access

No CISO function can be fully automated

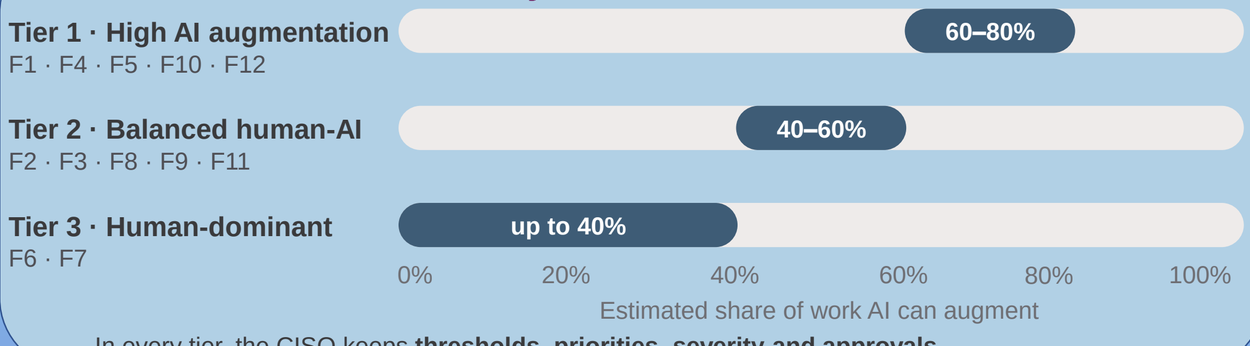

In every tier, the CISO keeps **thresholds, priorities, severity and approvals**.

Similar AI share, different oversight

Criteria	University	FinTech
AI augmentation	~72%	~74%
Human decision points	4	5
Regulatory density	Moderate	High (incl. DORA)
AI Act risk class	Limited risk	High risk (Annex III)
Processing	Batch	Real-time

The difference comes from **regulation, risk and real-time stakes**, not from what the AI can do.

Four phases: AI supports, humans decide

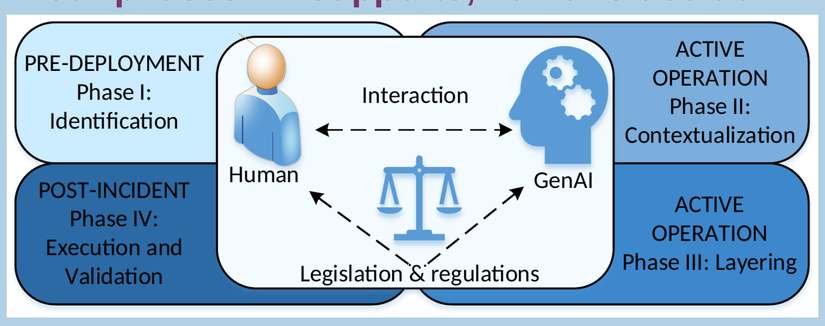

AI SUPPORTS
Policy mapping, data aggregation, scoring, monitoring, triage, reporting

HUMAN DECIDES
Risk thresholds, deployment, escalation, severity, recovery, notifications, sign-off

A human validates every action that is **irreversible, legally accountable or context-dependent**.

Key findings

1. **No CISO function can be fully automated** — Even data-heavy tasks need human judgment at key points
2. **Regulation creates compound accountability** — One decision can trigger obligations under several frameworks
3. **ISO 37000 provides the governance logic** — Though it does not define CISO- or AI-specific roles
4. **Sector context sets the level of oversight** — More rules and higher stakes mean more human decision points

Limitations: mapping based on ISO/IEC 27002:2022 · AI capabilities will change · illustrative use cases

Acknowledgments. The research was funded by the project “Research on Cyber Resilience Through Application of Generative Artificial Intelligence in Chief Information Security Officer Operations” (GAISO), which received funding from the Research Council of Lithuania (LMTLT) under agreement No. S-ITP-24-13.