Table of Contents

1 GRC & AI Governance Production Portfolio Binder

Prepared: 14 August 2026

Portfolio integrity statement: This binder demonstrates operating artifacts and public-source analysis. It does not claim senior employment tenure, certification, audit opinion, legal advice or client relationships that do not exist.

---

2 Technology GRC, TPRM & AI Governance Proof-of-Work Portfolio

GRC SOC2 ISO27001 NISTAIRMF EUAIAct TPRM

2.1 Purpose

This repository demonstrates operational Technology GRC, Third-Party Risk Management, customer assurance, privacy processor governance, AI governance, and audit-readiness work through ten completed proof-of-work systems.

The portfolio is intentionally evidence-first. It does not claim that a public-source case study is a certification, penetration test, legal opinion, SOC 2 examination, ISO audit, or substitute for live enterprise ownership. It demonstrates the analyst and operating artifacts used to perform those responsibilities.

2.2 Current 2026 framework status

- NIST AI RMF 1.0 remains the current published framework, and NIST states that it is being revised: https://www.nist.gov/itl/ai-risk-management-framework
- NIST's Generative AI Profile remains the current cross-sector GenAI companion profile: https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
- EU AI Act Article 50 transparency obligations apply from 2 August 2026; the Commission published implementation guidelines on 20 July 2026: https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems
- ISO/IEC 27001:2022 remains the current ISMS requirements standard: https://www.iso.org/standard/27001
- ISO/IEC 42001:2023 remains the AI management-system standard: https://www.iso.org/standard/42001

- The AICPA Trust Services Criteria remain the control criteria used for SOC 2 engagements: https://www.aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022

2.3 Portfolio projects

1. **Enterprise Trust & Customer Assurance Engine** - public-source procurement evidence analysis and customer assurance.
2. **AI Governance Operating System** - AI inventory, ownership, risk scoring, evaluation and monitoring.
3. **Enterprise TPRM & AI Subprocessor Engine** - vendor tiering, vendor evidence review and security questionnaire.
4. **SOC 2 + ISO 27001 Control-to-Evidence System** - unified control inventory, evidence and test procedures.
5. **Executive Technology Risk & Board Reporting System** - risk register, KRIs, appetite and escalation.
6. **AI Transparency & Article 50 Audit System** - interactive/generative AI transparency and content provenance.
7. **Shadow AI & Prompt DLP Governance System** - data classification, local redaction and AI egress incident response.
8. **Enterprise Procurement & Security Triage System** - 25 high-friction buyer questions with evidence-backed answers.
9. **GDPR Article 28 DPA & Subprocessor Governance Engine** - processor clauses, subprocessor authorization and audit evidence.
10. **Continuous Control Monitoring & Audit Operations System** - evidence calendar, audit request tracker and exception workflow.

2.4 Artifact map

```
GRC_AI_Governance_Portfolio/
├── 00_MASTER_README.md
├── 01_Enterprise_Trust_Customer_Assurance.md
├── 02_AI_Governance_Operating_System.md
├── 03_TPRM_AI_Subprocessor_Engine.md
├── 04_SOC2_ISO27001_Control_Evidence.md
├── 05_Executive_Technology_Risk.md
├── 06_AI_Transparency_Article50.md
├── 07_Shadow_AI_Prompt_DLP.md
├── 08_Procurement_Security_Triage.md
├── 09_GDPR_Article28_DPA_Engine.md
├── 10_Continuous_Audit_Operations.md
├── 11_LinkedIn_Deployment_Suite.md
├── 12_Personal_Website_Copy.md
├── 13_Global_Contractor_Positioning.md
└── 14_Source_Register.md
```

2.5 Companion workbook

The companion spreadsheet GRC_AI_Governance_10_Project_Evidence_Workbook.xlsx contains the populated registers and operational tables referenced by these project documents.

2.6 Evidence integrity rule

Every external claim in this portfolio is classified as one of the following:

- **Observed public evidence** - directly supported by a cited source.
- **Not publicly identified** - the reviewed public material did not expose the evidence; this never means the control does not exist internally.
- **Reference enterprise implementation** - a complete simulated operating environment built to demonstrate the workflow without pretending it is a real client's confidential program.

This distinction is mandatory throughout the portfolio.

3 Project 1 - Enterprise Trust & Customer Assurance Engine

3.1 Executive objective

Demonstrate the ability to turn scattered public security, privacy and AI statements into a defensible enterprise procurement evidence model. The project uses Lexi, a live legal-AI SaaS product, as a public-source case study because legal data carries unusually high confidentiality and buyer-assurance expectations.

Scope boundary: public-source readiness review only. No claim is made about non-public controls, actual SOC 2 findings, legal compliance, vulnerability status, or attorney-client privilege outcomes.

3.2 Verified public strengths

1. Lexi publishes dedicated Terms, Privacy and Data Controls pages and explicitly names OpenAI and Anthropic as third-party AI providers. Source: https://www.getlexi.io/data-controls
2. Lexi states that customer work is account-specific and not used to train models for other users. Source: https://www.getlexi.io/data-controls
3. Lexi publicly describes access controls, audit trails, encryption and customer deletion rights. Source: https://www.getlexi.io/data-controls

3.3 Three procurement-friction findings

3.3.1 Finding 1 - AI provider retention wording

Lexi's Data Controls describes zero-data-retention arrangements and says third-party generative-AI providers do not retain customer data beyond what is required for the immediate request. The June 2026 Terms separately state that data shared with OpenAI and Anthropic is

deleted after 30 days. Both statements may be reconcilable through endpoint, provider, logging or contract differences, but the public wording requires a buyer to reconcile them manually.

Buyer question: Which production routes operate under zero-data-retention terms, and which routes may retain content or metadata for up to 30 days?

Required evidence: provider-by-provider retention matrix, exact product/endpoints covered, abuse-monitoring exceptions, regional processing, contractual record and current date.

Sources: https://www.getlexi.io/data-controls and https://www.getlexi.io/terms

3.3.2 Finding 2 - customer-content training language

Lexi's Data Controls emphasizes that legal work is not used to train models for other users. The Terms state that customers retain control over whether content may be used to train Lexi AI models through a service setting.

Buyer question: Is training impossible, default-off, explicit opt-in, limited to tenant-specific adaptation, or available for some accounts?

Required evidence: one authoritative taxonomy covering inference, retrieval, evaluation, product improvement, tenant-specific personalization, fine-tuning and generalized model training.

Sources: https://www.getlexi.io/data-controls and https://www.getlexi.io/terms

3.3.3 Finding 3 - retention schedule specificity

The Data Controls page directs customers to review specific retention periods in the Privacy Policy, while the current Privacy Policy primarily describes retention using necessity and purpose criteria rather than a simple object-level product schedule.

Buyer question: What are the default deletion windows for documents, prompts/outputs, logs, analytics, backups and closed accounts?

Required evidence: product data retention table and deletion/offboarding SLA.

Sources: https://www.getlexi.io/data-controls and https://www.getlexi.io/privacy

3.4 15-point evidence architecture

The companion workbook P01_Trust_Readiness contains the complete evidence matrix. The operating schema is:

Field	Purpose
Area	Security/privacy/AI topic
Buyer Question	What an enterprise reviewer needs answered
Finding Type	Observed / Not publicly identified / Internal verification required
Required Evidence	Artifact needed to answer the buyer

Field	Purpose
	defensibly
Framework	SOC 2 / ISO 27001 / GDPR / NIST / Article 50 mapping
Priority	High / Medium / Low
Owner	Accountable function
Status	Open / In progress / Closed

3.5 Customer assurance knowledge-base rule

Every canonical customer-assurance answer must include:

1. one accountable owner;
2. one or more evidence sources;
3. a last-reviewed date;
4. a review-expiry date;
5. an accuracy classification;
6. a statement of scope;
7. explicit prohibition on claiming certification, encryption, retention, data residency, model-training behavior or legal compliance without current evidence.

3.6 Executive remediation roadmap

0-48 hours

- Normalize AI-provider retention language across Terms, Privacy, Data Controls and sales questionnaire responses.
- Normalize the customer-content training statement and define the default setting plus allowed learning modes.
- Publish a concise retention table covering core customer content, AI prompts/outputs, logs, backups and account deletion.

Week 1

- Build a named model-provider/subprocessor evidence register.
- Establish a controlled customer assurance library with owners and review dates.
- Define an assurance language standard distinguishing certification, attestation, alignment and internal design claims.

Ongoing

- Trigger an out-of-cycle review whenever a model provider, retention setting, region, subprocessor, assurance report or major product architecture changes.

3.7 Executive pitch

I convert a SaaS company's public trust claims into a procurement-ready evidence system that shows exactly what a buyer can verify, what remains ambiguous, and what documentation should be fixed before sales or engineering repeats the same answers across enterprise deals.

4 Project 2 - AI Governance Operating System

4.1 Objective

Create a complete AI governance operating model for a reference global B2B SaaS organization using current NIST AI RMF, the NIST Generative AI Profile and ISO/IEC 42001. The system is designed to prove governance work at the operating-artifact level: inventory, risk ownership, evaluation, monitoring, escalation and retirement.

NIST AI RMF 1.0 remains the current published framework but is under revision. Its core functions are GOVERN, MAP, MEASURE and MANAGE. Source: https://www.nist.gov/itl/ai-risk-management-framework and https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

ISO/IEC 42001:2023 specifies requirements for an AI management system. Source: https://www.iso.org/standard/42001

4.2 AI governance policy

4.2.1 1. Purpose

The organization shall identify, govern, assess, test, approve, monitor and retire AI systems in proportion to their impact on customers, employees, business decisions, data confidentiality, legal obligations and service availability.

4.2.2 2. Scope

This policy applies to externally sourced AI services, internally developed models, embedded AI features, AI agents, generative AI, machine-learning classifiers, AI-assisted decision systems, foundation-model APIs and automated content-generation tools used for organizational business.

4.2.3 3. Inventory requirement

No production AI system may operate without an inventory record containing business purpose, owner, technical owner, provider/model class, data categories, affected stakeholders, human oversight, deployment region, third-party dependencies, regulatory exposure, intended outcome, material limitations and current approval status.

4.2.4 4. Risk ownership

Every AI system shall have one accountable business owner and one technical owner. High-impact use cases shall additionally have a named risk or compliance reviewer. Owners are responsible for keeping inventory and risk evidence current.

4.2.5 5. Pre-deployment governance

Before production use, the owner shall document intended purpose, foreseeable misuse, affected stakeholders, data sensitivity, autonomy level, human-review points, third-party/model risks, failure modes and measurable acceptance criteria.

4.2.6 6. Evaluation

AI systems shall be evaluated using task-appropriate metrics. Evaluation may include factual accuracy, groundedness, toxicity, privacy leakage, prompt-injection resistance, refusal behavior, bias/fairness, security, latency, drift, failure recovery and human-review effectiveness. A system shall not be described as safe or compliant merely because a generic benchmark score is high.

4.2.7 7. Human oversight

AI may assist material decisions, but material legal, financial, security, employment, regulatory or customer-impacting decisions require a qualified human decision owner unless an approved risk assessment explicitly permits automation under applicable law and organizational policy.

4.2.8 8. Data governance

AI data use shall distinguish inference, retrieval, evaluation, product improvement, tenant-specific personalization, fine-tuning and generalized training. Customer or confidential data may only be sent to an AI provider approved for the relevant classification and contract terms.

4.2.9 9. Change management

A material change in model, provider, prompt gateway, retrieval source, autonomy, data class, geography, intended purpose or safety control requires governance review before production rollout.

4.2.10 10. Monitoring

High and medium risk AI systems shall have KRIs or control metrics. Examples include hallucination severity, policy-refusal failures, PII leakage attempts, unsupported-answer rate, human-override rate, drift, incident volume and unresolved high-risk findings.

4.2.11 11. Incident escalation

AI incidents involving data leakage, deceptive content, unauthorized material action, discriminatory outcome, high-severity incorrect advice, regulatory exposure or systematic control failure shall be routed into the enterprise incident-management process.

4.2.12 12. Retirement

Retired AI systems shall be removed from production access, credentials revoked, data handling closed, residual data addressed under retention policy, and the inventory record marked retired with effective date and evidence.

4.3 Completed AI system inventory

The companion workbook P02_AI_Governance contains fifteen fully populated AI use cases including employee assistant, customer chatbot, support copilot, security questionnaire assistant, contract review, vendor risk summarization, code assistant, meeting summarization, marketing generation, HR support, finance analytics, fraud prioritization, classification and executive risk summarization.

4.4 Risk scoring

Inherent risk = Likelihood (1-5) x Impact (1-5).

Residual risk is assigned only after documenting the controls that reduce likelihood or impact. A residual score is never reduced solely because a vendor advertises enterprise security.

4.5 Required AI evidence pack

- AI System Inventory
- AI Risk Register
- Model/Provider Register
- Data-Use & Training Matrix
- Evaluation Plan and Results
- Human Oversight Matrix
- Article 50 Applicability Record for EU-facing systems
- AI Change Register
- AI Incident Register
- Quarterly AI Governance Review Minutes

4.6 Executive pitch

> I operationalize AI governance as an accountable system: every AI use case has a purpose, owner, data boundary, risk decision, evaluation evidence, human-oversight rule and monitoring signal instead of relying on a high-level responsible-AI policy.

5 Project 3 - Enterprise TPRM & AI Subprocessor Engine

5.1 Objective

Create a risk-tiered third-party risk management system for a B2B SaaS company using real current enterprise technology vendors. The companion workbook contains ten vendor records and a twenty-question vendor security questionnaire.

GDPR/UK GDPR Article 28 processor governance requires appropriate contracts and controls around subprocessors; ICO guidance describes documented instructions, confidentiality, security, subprocessor authorization, rights assistance, deletion and audit/information obligations. Source: https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/contracts-and-liabilities-between-controllers-and-processors-multi/what-needs-to-be-included-in-the-contract/

5.2 Vendor tiering method

5.2.1 Tier 1 - Critical

A vendor is Tier 1 if outage, breach or control failure could materially affect production availability, source code, privileged access, payment flows, regulated/customer data or core AI processing.

5.2.2 Tier 2 - High

A vendor is Tier 2 if it receives sensitive telemetry or business data or materially supports operations but does not independently control a mission-critical production path.

5.2.3 Tier 3 - Moderate

A vendor is Tier 3 if it processes internal business data with limited production dependency.

5.2.4 Tier 4 - Low

A vendor is Tier 4 if it processes only public/non-sensitive data and has low operational dependency.

5.3 Public-source vendor examples

5.3.1 OpenAI

OpenAI states that business/API data is not used to train models by default and documents API retention controls, including 30-day removal for typical API data and zero-data-retention eligibility for qualifying use cases. Source: https://openai.com/enterprise-privacy/

5.3.2 Amazon Bedrock

AWS states that Bedrock protects data with encryption in transit/at rest, does not share customer inputs/outputs with model providers and does not use them to train base foundation models. Source: https://aws.amazon.com/bedrock/security-privacy-responsible-ai/

5.3.3 Gemini Enterprise

Google documents enterprise controls including encryption, supported compliance certifications, CMEK, VPC Service Controls, Access Transparency and data-residency capabilities depending on edition and region. Source: https://docs.cloud.google.com/gemini/enterprise/docs/compliance-security-controls

5.3.4 Slack

Slack documents encryption at rest/in transit, EKM, audit logs, native DLP, data residency and ISO/SOC assurance. Source: https://slack.com/trust/security

5.3.5 GitHub Enterprise Cloud

GitHub documents regional storage for Enterprise Cloud data residency and identifies categories that may remain or be processed outside the selected region depending on function. Source: https://docs.github.com/en/enterprise-cloud@latest/admin/data-residency/about-storage-of-your-data-with-data-residency

5.4 Twenty-question vendor questionnaire

The complete questionnaire is contained in workbook sheet P03_Vendor_Questions. Questions cover governance, assurance, encryption, IAM, tenant isolation, subprocessors, retention, deletion, model training, provider retention, secure SDLC, vulnerabilities, incident response, resilience, logging, privacy rights, transfers, AI governance and risk exceptions.

5.5 Vendor decision standard

Each vendor review must end in exactly one of four states:

1. **Approve** - residual risk is within appetite.
2. **Approve with conditions** - use is permitted only after named controls or contract changes.
3. **Escalate for risk acceptance** - business wishes to proceed above appetite; designated risk owner must approve time-bound acceptance.
4. **Reject** - risk is unacceptable or required evidence/terms are unavailable.

5.6 TPRM workflow

Request -> Inherent Risk -> Tier -> Evidence -> Analyst Review -> Findings -> Remediation/Contract -> Decision -> Monitoring -> Renewal

5.7 Subprocessor disclosure clause

> The Processor shall not appoint a new or replacement Subprocessor to process Controller Personal Data without the Controller's prior specific or general written authorization. Where general authorization applies, the Processor shall provide advance written notice identifying the proposed Subprocessor, the processing purpose, location and effective date, and shall provide a reasonable period for the Controller to object on legitimate data-protection grounds. The Processor shall impose written data-protection obligations on each Subprocessor that provide an equivalent level of protection for the obligations applicable to the Processor under the parties' data-processing agreement. The Processor remains responsible to the Controller for the Subprocessor's performance of those obligations.

This clause is a portfolio drafting example aligned to Article 28 concepts and is not legal advice or a substitute for counsel review.

5.8 Executive pitch

I turn vendor onboarding into a risk decision by tiering critical vendors, validating evidence, tracing AI subprocessors and data flows, documenting residual risk and producing a defensible approve-remediate-accept-or-reject outcome.

6 Project 4 - SOC 2 + ISO 27001 Control-to-Evidence System

6.1 Objective

Demonstrate the operational layer between a framework and an audit: control narratives, owners, recurring evidence, tests, exceptions and remediation.

The AICPA Trust Services Criteria provide criteria for Security, Availability, Processing Integrity, Confidentiality and Privacy used in SOC 2 engagements. Source: https://www.aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022

ISO/IEC 27001:2022 defines requirements for establishing, implementing, maintaining and continually improving an information security management system. Source: https://www.iso.org/standard/27001

6.2 Control design standard

Every production control record must contain:

- unique control ID;
- control objective;
- precise control statement;
- mapped framework criteria;
- accountable owner;
- evidence source;
- operating frequency;
- preventive/detective and manual/automated classification where relevant;
- test procedure;
- exception process;
- current status.

6.3 Completed control domains

The companion workbook P04_Control_Evidence contains fifteen domains:

1. Security Governance
2. Risk Management
3. Asset Management

4. Identity & Access
5. Privileged Access
6. Encryption
7. Logging & Monitoring
8. Vulnerability Management
9. Secure SDLC / Change Management
10. Incident Response
11. Business Continuity / Disaster Recovery
12. Vendor Risk
13. Privacy / Processor Governance
14. AI Governance
15. Customer Assurance

6.4 Continuous evidence collection schedule

Evidence	Frequency	Owner	Quality Gate
Governance meeting minutes	Quarterly	GRC	Approval and decisions recorded
Enterprise risk register export	Quarterly	GRC	Owners, treatments and overdue items current
User access review	Quarterly	IT	Reviewer sign-off and remediation closed
Privileged access review	Monthly	Security	All privileged accounts attributable and justified
SIEM coverage report	Monthly	Security	All critical systems present and retained
Vulnerability SLA dashboard	Monthly	Security	Critical/high overdue items escalated
Secure release sample	Monthly	Engineering	Code review and required tests present
IR tabletop / exercise	Annual	Security	Findings and actions tracked to closure
Restore/failover test	Semiannual	Engineering	RTO/RPO evidence recorded
Tier 1 vendor review	Annual / renewal	GRC	Evidence current and findings resolved/accepted
Subprocessor register review	Quarterly	Privacy	New vendors and notices reconciled

Evidence	Frequency	Owner	Quality Gate
AI system risk review	Quarterly / release	AI Governance	Material changes reviewed
Customer assurance library	Quarterly	GRC	Answers remain tied to current evidence

6.5 Control testing methodology

A test conclusion must distinguish:

- **Design effective** - the control as documented could reasonably meet its objective.
- **Operating effective** - sampled evidence shows the control operated as designed during the test period.
- **Exception** - evidence does not meet the control requirement.
- **Not tested** - insufficient population, period or evidence.

No portfolio project may claim a SOC 2 Type II conclusion because Type II is an independent attestation over control operation for a defined period.

6.6 Executive pitch

> I translate SOC 2 and ISO 27001 into operating controls with owners, evidence schedules and test procedures, so an audit request points to an existing control system instead of starting a document hunt.

7 Project 5 - Executive Technology Risk & Board Reporting System

7.1 Objective

Convert technical GRC activity into management decisions using a centralized enterprise risk register, risk appetite, KRIs, escalation thresholds and treatment ownership.

7.2 Risk statement standard

Risk statements use the format:

> Because of **cause / threat**, there is a possibility that **risk event** will occur, resulting in **business impact**.

Risks are never written as vague topics such as “cybersecurity risk” or “AI risk.”

7.3 Risk scoring

Inherent risk = Likelihood (1-5) x Impact (1-5).

Residual risk is the risk remaining after documented controls. A lower residual score must be justified by control effectiveness rather than subjective optimism.

7.4 Completed risk universe

The workbook P05_Executive_Risk contains fifteen risks covering SaaS outage, privileged compromise, shadow-AI leakage, subprocessor breach, Article 50 transparency, customer assurance overstatement, critical vulnerabilities, deletion failure, AI quality harm, logging gaps, audit evidence lateness, code/secret leakage, processor-contract gaps, stale trust claims and AI provider/model changes.

7.5 Board reporting cadence

7.5.1 Monthly operational risk review

- New high/critical risks
- Risks that crossed appetite
- Overdue risk treatments
- Security and AI incidents
- Audit/control exceptions
- Critical vendor issues

7.5.2 Quarterly executive pack

1. Top five enterprise technology risks.
2. Risk movement since prior quarter.
3. Risks outside appetite.
4. Top KRIs and threshold breaches.
5. Audit and customer-assurance posture.
6. AI governance posture.
7. Third-party concentration and critical findings.
8. Decisions required from leadership.

7.6 KRI library

- Percentage of Tier 1 vendors overdue for assessment.
- Percentage of critical controls missing current evidence.
- Critical vulnerabilities past SLA.
- Privileged accounts without required MFA.
- High-risk AI systems lacking accountable owner.
- EU-facing AI systems lacking Article 50 applicability assessment.
- Customer assurance answers lacking current evidence.
- Shadow-AI events involving confidential data.
- Data deletion requests exceeding SLA.
- Audit exceptions overdue for remediation.
- Production AI provider/model changes without governance review.
- Critical systems missing required security logs.

7.7 Executive pitch

I convert security and compliance activity into a decision system: what can materially go wrong, how much risk remains, whether it is inside appetite, who owns treatment and what leadership must decide.

8 Project 6 - Multi-Framework AI Governance & Article 50 Transparency Audit System

8.1 Objective

Operationalize Article 50 transparency for a reference enterprise AI portfolio while mapping governance to NIST AI RMF and ISO/IEC 42001.

The European Commission states that Article 50 transparency obligations apply from 2 August 2026 and cover certain interactive/generative AI systems, machine-readable marking of AI-generated/manipulated content, and deployer disclosures for specified deepfake, emotion-recognition, biometric-categorization and public-interest content scenarios. Source: https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems

8.2 Completed 15-point transparency register

The companion workbook P06_AI_Transparency contains fifteen use cases, including customer chatbots, voice agents, synthetic images/video, executive voice clones, public-interest publications, recruiting chatbots, emotion recognition, biometric categorization, internal meeting summarization, AI-generated knowledge articles, product mockups, investor communications and translation/dubbing.

Each row records:

- provider/deployer role;
- direct human interaction;
- synthetic content type;
- public-interest text relevance;
- emotion/biometric use;
- potential Article 50 trigger;
- required transparency control;
- machine-readable marking/provenance requirement;
- NIST AI RMF anchor;
- risk and accountable owner.

9 Synthetic Content & Deepfake Disclosure Control Policy

9.1 1. Purpose

This policy establishes mandatory transparency, provenance, approval and labelling controls for AI-generated or AI-manipulated text, image, audio and video produced or deployed by the organization.

9.2 2. Scope

The policy applies to marketing, communications, product features, customer service, recruiting, investor communication, public-interest publications, executive likenesses, synthetic personas, dubbing/translation, content moderation and any other organizational use of generative or manipulative AI.

9.3 3. Core rule

AI-generated or manipulated content must not be presented in a way that is materially deceptive about its origin, authorship, authenticity or the identity/actions of a real person. Where Article 50 or other applicable requirements impose disclosure or marking, the designated product/legal owner shall implement and evidence the requirement before release.

9.4 4. Interactive AI

Customer-facing interactive AI shall provide a clear AI-interaction disclosure unless the AI nature is obvious from the circumstances and the applicable legal assessment confirms no additional disclosure is required. High-risk contexts shall use explicit disclosure regardless of ambiguity.

9.5 5. Synthetic media provenance

For generated or materially manipulated image, audio and video, the organization shall preserve available provenance metadata, generation records and content-review evidence. Where technically feasible and legally applicable, machine-readable marking shall be enabled and preserved through the publication workflow.

9.6 6. Deepfakes and real-person likeness

Synthetic content depicting a real identifiable person saying or doing something they did not actually say or do requires Legal/Communications approval before external release. Required deepfake disclosure shall be conspicuous enough for the intended audience to understand that the content is artificial or manipulated.

9.7 7. Public-interest text

AI-generated or manipulated text published to inform the public on matters of public interest shall receive documented human editorial review before publication. If the organization intentionally publishes without the human review/editorial responsibility contemplated by applicable law, Legal shall determine and implement the required Article 50 disclosure.

9.8 8. Emotion recognition and biometric categorization

No emotion-recognition or biometric-categorization system may be deployed without documented legal review. Where deployment is lawful, required notification to affected persons shall occur before or at exposure to the system.

9.9 9. Marketing controls

Marketing may use generative AI only under approved brand and provenance controls. AI output shall not falsely imply product capability, customer endorsement, real-world event evidence or a real-person statement.

9.10 10. Records

The owner shall retain the use-case decision, disclosure text, marking/provenance evidence, human approval and legal assessment for each high-risk or externally distributed synthetic-content workflow.

9.11 11. Monitoring

Quarterly review shall test a sample of externally distributed AI content for required disclosure, retained provenance and policy compliance. Any systematic failure is escalated as an AI governance issue.

9.12 12. Exceptions

Policy exceptions require written business justification, legal/risk review, compensating controls, approver and expiration date. There are no permanent exceptions.

9.13 CTO remediation roadmap

9.13.1 0-48 hours

- Inventory all customer-facing generative/interactive AI features.
- Add provider/deployer role and Article 50 applicability fields to the release checklist.
- Add a standard AI interaction disclosure component to customer-facing AI UI.
- Preserve provenance metadata for new synthetic media assets.
- Block synthetic real-person likeness release without Legal/Communications approval.

9.13.2 30 days

- Implement a content-provenance standard.
- Create automated release-gate evidence for required disclosures.
- Add Article 50 controls to AI change management and product QA.

9.14 Executive pitch

> I translate Article 50 from legal text into a release-control system that identifies which AI interactions or synthetic content need disclosure, marking, human review and retained evidence before they reach users.

10 Project 7 - Enterprise Shadow AI & Local Prompt Redaction Governance

10.1 Objective

Reduce the probability that employees disclose confidential, regulated or credential data to unsanctioned AI tools while preserving legitimate enterprise AI productivity.

The NIST Generative AI Profile provides cross-sector risk-management actions for generative AI. Source: https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

11 Local Memory Vault & AI Data Retention Policy Specification

11.1 1. Approved storage principle

The organization's approved AI workflows shall minimize external persistence of confidential prompts and retrieved content. Where a local or organization-controlled memory layer is required, it shall use encrypted storage, explicit retention, least-privilege access and auditable deletion.

11.2 2. Local memory vault

The local memory vault is an organization-controlled data store used for temporary AI context, approved retrieval indexes, tokenized identifiers and interaction metadata. It must not become an unmanaged duplicate system of record.

11.3 3. Data minimization

Only the minimum data required for the approved AI task may be stored. Credentials, authentication secrets, private keys, raw payment-card data and prohibited regulated data must not be stored in AI memory.

11.4 4. Retention classes

- **Transient prompt buffer:** delete after the active workflow unless a documented business requirement exists.
- **AI interaction logs:** 30 days by default for low-risk troubleshooting; extend only with approved security/legal requirement.
- **Evaluation samples:** 90 days unless anonymized/de-identified and separately approved.
- **Tenant retrieval index:** retained while the source remains active and authorized; deleted within 30 days of source removal or account closure unless a shorter contract applies.
- **Security incident evidence:** retained under the incident/evidence retention schedule.
- **Legal-hold data:** retained only under documented legal hold.

11.5 5. Encryption

Vault data shall be encrypted at rest and in transit. Production key material shall be managed through the organization's approved key-management service. Application code shall not contain static production keys.

11.6 6. Access

Access is least-privilege and role-based. Administrative access is logged. Production vault access by engineering or support requires approved need and is reviewed periodically.

11.7 7. Redaction

Before confidential data is submitted to an external AI provider, approved gateways shall apply deterministic and/or classifier-based redaction for secrets, credentials, high-risk identifiers and prohibited regulated data. Redaction shall preserve the minimum context required for the task.

11.8 8. Deletion

Deletion shall remove the logical object from the active store and follow documented backup purge behavior. Deletion failures or orphaned vector/index content are incidents when contractual or privacy obligations are missed.

12 Data Classification Standard for AI Use

Class	Examples	AI Rule
Public	Approved website copy, public product docs	May use approved AI services
Internal	Internal procedures, non-sensitive plans	Approved enterprise AI only
Confidential	Customer contracts, source code, private financials, non-public strategy	Approved enterprise AI plus use-case approval; DLP/redaction required where applicable
Restricted	Credentials, secrets, raw payment-card data, highly regulated identifiers, privileged legal material without specific approval	Prohibited from general AI prompts; only purpose-built approved workflow with documented legal/security authorization

13 Data Flow Architecture

User / Application
    |
    v

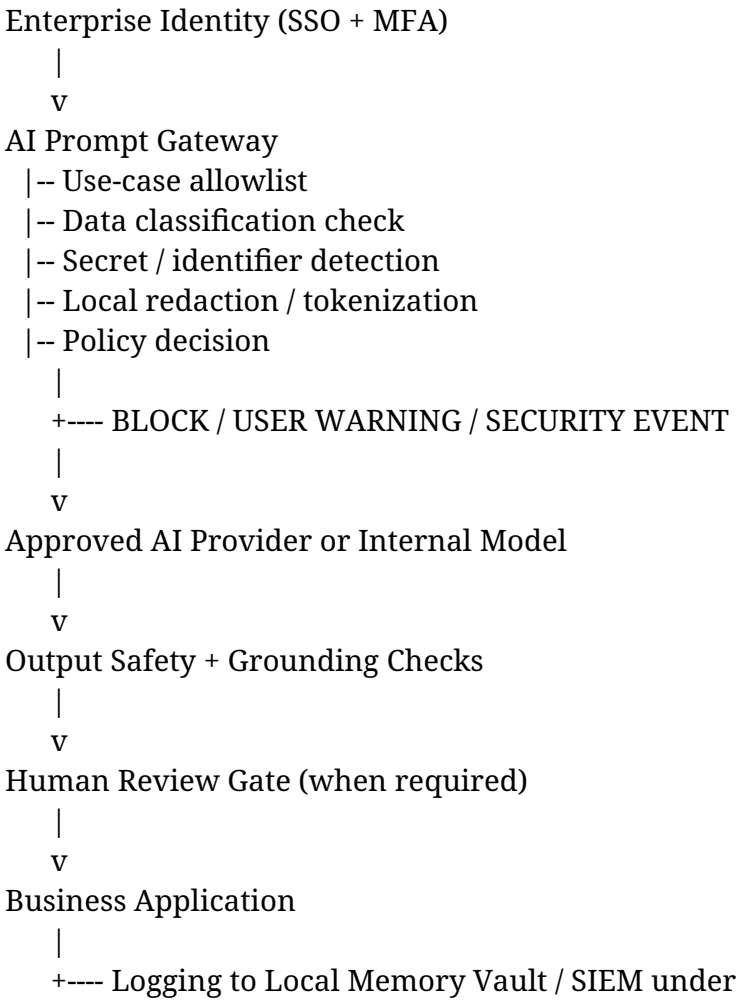

14 Incident Response Playbook - Unsanctioned AI Egress

14.1 Trigger

Initiate this playbook when DLP, CASB, endpoint monitoring, employee report or audit identifies that confidential/restricted information was submitted to an unapproved AI service or an approved service outside its authorized use case.

14.2 Severity

- **SEV-1:** credentials, private keys, regulated high-impact data, material customer secrets, privileged material, or confirmed external exposure.
- **SEV-2:** confidential data submitted to unapproved provider with uncertain retention/access.
- **SEV-3:** internal data submitted contrary to policy without high sensitivity.

14.3 First 15 minutes

1. Preserve event evidence: user, time, tool, account, data class, URL/app, prompt size and available content fingerprint.
2. Invalidate exposed credentials immediately if secrets are involved.
3. Block or quarantine the unapproved AI application/account where technically possible.
4. Notify Security incident lead and Privacy/Legal for SEV-1/SEV-2.

14.4 First hour

1. Determine data subject/customer/system scope.
2. Determine whether the provider account is personal or enterprise and identify provider retention/deletion options.
3. Request deletion through provider controls/support when available.
4. Determine whether the data was used for model training or product improvement under the provider/account terms.
5. Identify downstream copies such as browser history, local exports, synced notes or screenshots.
6. Assess contractual and regulatory notification obligations.

14.5 First 24 hours

1. Complete root-cause interview and system evidence collection.
2. Document risk to confidentiality, privacy, privilege, intellectual property and credentials.
3. Confirm containment/deletion evidence.
4. Notify affected customers/regulators only under approved Legal/Privacy determination.
5. Add DLP/CASB detection if the event exposed a new pattern.

14.6 Seven-day closure

1. Complete root cause and lessons learned.
2. Correct policy, allowlists, DLP or training gaps.
3. Review whether the same workflow exists elsewhere.
4. Record final severity, exposure conclusion and residual risk.
5. Track corrective actions to closure.

14.7 Executive pitch

> I govern shadow AI as a data-loss problem with an allowlist, local redaction, classification-aware prompt gateway, measurable monitoring and a real incident playbook instead of relying on employee awareness alone.

15 Project 8 - Automated Enterprise Procurement & Security Triage System

15.1 Objective

Turn repetitive 100-question enterprise security questionnaires into a controlled knowledge base with approved, evidence-backed answers. The system does not auto-approve legal or security claims: it accelerates drafting while preserving accountable review.

15.2 Master Security & Privacy Knowledge Base

15.2.1 1. Do you maintain a formal information security program?

Standardized response

Yes. The reference control model defines accountable security governance, approved policies, management review, enterprise risk management, control ownership and recurring evidence review. Production claims must always be supported by current approved records.

Evidence required: Security Governance Policy; governance minutes; org chart.

15.2.2 2. Are you SOC 2 Type II certified/compliant?

Standardized response

Use exact assurance language only. State current report type, period and scope if an examination exists; otherwise describe designed/aligned controls without implying an issued SOC 2 report.

Evidence required: Current SOC 2 report or audit roadmap.

15.2.3 3. Are you ISO 27001 certified?

Standardized response

If certified, provide certificate scope, certification body and validity. If not, state alignment or implementation status without using certified.

Evidence required: ISO certificate or ISMS roadmap.

15.2.4 4. How is customer data encrypted?

Standardized response

Customer data is encrypted in transit using approved TLS and at rest using platform/cloud encryption. Key ownership, exceptions and customer-managed-key options are documented per system.

Evidence required: Architecture; KMS configuration; cryptography standard.

15.2.5 5. How do you control privileged access?

Standardized response

Privileged access is individually attributable, least-privilege, MFA-protected, approved, logged and periodically reviewed.

Evidence required: PAM/admin roster; access review; MFA evidence.

15.2.6 6. Do you support SSO and MFA?

Standardized response

Enterprise authentication uses SSO where supported and MFA is mandatory for privileged and remote administrative access. Product capability is stated per plan.

Evidence required: IdP configuration; product documentation.

15.2.7 7. How are vulnerabilities managed?

Standardized response

Vulnerabilities are identified through automated scanning and testing, risk-rated, assigned remediation SLAs, tracked to closure and escalated when overdue; exceptions require approval.

Evidence required: Vulnerability policy; scanner dashboard; tickets.

15.2.8 8. Do you perform penetration tests?

Standardized response

Independent penetration testing is performed on a risk-based cadence and after major changes where warranted. Only current approved summaries are shared externally.

Evidence required: Pen-test executive summary.

15.2.9 9. How do you respond to incidents?

Standardized response

An approved incident process defines severity, roles, triage, containment, investigation, evidence handling, communications and lessons learned; notification follows contract and law.

Evidence required: IR plan; exercise report.

15.2.10 10. What are your disaster-recovery controls?

Standardized response

Critical systems have recovery objectives, encrypted backups and tested restore/failover procedures with periodic recovery evidence.

Evidence required: BCP/DR; restore test; RTO/RPO.

15.2.11 11. Which subprocessors process customer data?

Standardized response

A maintained register identifies provider, service, location and purpose. New subprocessors follow the authorization/change process in the applicable DPA.

Evidence required: Subprocessor list; DPA.

15.2.12 12. How do you notify customers of subprocessor changes?

Standardized response

Under general authorization, customers receive advance notice of intended additions/replacements and an opportunity to object under the applicable DPA process.

Evidence required: DPA clause; notice procedure.

15.2.13 13. Do you use customer data to train AI models?

Standardized response

The approved AI data-use statement distinguishes inference, evaluation, tenant-specific personalization, product improvement and generalized training. The answer states the actual default and opt-in configuration.

Evidence required: AI data-use policy; provider terms.

15.2.14 14. Which AI providers receive customer data?

Standardized response

The AI provider register identifies each model/provider, purpose, data categories, region, retention, training behavior and approved use cases.

Evidence required: AI provider register; TPRM evidence.

15.2.15 15. How long do you retain customer data?

Standardized response

Retention is defined by data object and purpose. The schedule covers customer content, prompts/outputs, logs, backups, billing records and legal-hold exceptions.

Evidence required: Retention schedule.

15.2.16 16. Can customers delete or export data?

Standardized response

The offboarding process supports committed export/return and deletion, including documented backup purge timing and legal-retention exceptions.

Evidence required: Deletion/offboarding SOP.

15.2.17 17. How do you support data-subject rights?

Standardized response

Where acting as processor, the organization supports the controller through documented technical and organizational measures and routes requests under the DPA and privacy process.

Evidence required: DPA; DSR procedure.

15.2.18 18. Where is customer data hosted?

Standardized response

Hosting/processing regions are documented per product and deployment. Residency statements distinguish primary content from telemetry, support data and subprocessors.

Evidence required: Region matrix; architecture.

15.2.19 19. How do you monitor third-party risk?

Standardized response

Vendors are tiered by data access, criticality and regulatory impact before onboarding and reassessed on a risk-based cadence. Findings are tracked to treatment or approved acceptance.

Evidence required: TPRM policy; vendor register.

15.2.20 20. How do you manage secure software changes?

Standardized response

Code changes require review, automated security testing and controlled deployment; secrets, dependency and vulnerability scanning operate in CI/CD.

Evidence required: Secure SDLC; CI evidence.

15.2.21 21. Do you log customer/admin activity?

Standardized response

Security-relevant events are centrally logged with defined retention and access controls. Customer-facing audit-log capability is stated separately from internal security logging.

Evidence required: Logging standard; SIEM coverage.

15.2.22 22. How do you govern AI risk?

Standardized response

AI systems are inventoried, assigned accountable owners, assessed for intended use and affected stakeholders, evaluated before/after release and monitored with escalation and retirement criteria.

Evidence required: AI policy; inventory; risk register.

15.2.23 23. How do you address EU AI Act Article 50?

Standardized response

Each EU-facing interactive/generative use case is assessed for provider/deployer role and specific duties. Applicable disclosure, labelling/marking and evidence are implemented before release.

Evidence required: Article 50 register; UI/provenance evidence.

15.2.24 24. How do you handle security exceptions?

Standardized response

Exceptions require business justification, risk assessment, compensating controls, expiry date and designated risk-owner approval; expired exceptions are escalated.

Evidence required: Exception register.

15.2.25 25. How do you keep questionnaire answers accurate?

Standardized response

A controlled customer-assurance library assigns every answer an owner, evidence source and review/expiry date. Material architecture, vendor or assurance changes trigger out-of-cycle review.

Evidence required: Answer library; change log.

15.3 Triage automation logic

1. Normalize the incoming questionnaire into canonical questions.
2. Match each item to the approved knowledge base.
3. Auto-fill only answers classified as current and evidence-backed.
4. Route legal/privacy questions to Privacy/Legal when the answer depends on contract, jurisdiction or role.
5. Route architecture/security questions to Engineering/Security when evidence is stale or implementation-specific.
6. Route AI/model questions to AI Governance/TPRM when provider, retention or training behavior has changed.
7. Block submission when a high-risk answer has no current evidence.
8. Save the final response, source evidence and reviewer approval for reuse.

15.4 Executive Trust Readiness Snapshot Summary

Public/Customer Assurance Status: Moderate - strong control structure, with priority on evidence freshness and provider-specific AI transparency.

Three strengths:

- Unified control/evidence library covering governance, IAM, encryption, logging, vulnerability, SDLC, incident, resilience, vendor risk, privacy and AI.
- Central third-party/model-provider register with risk tiering.
- Article 50 and AI governance integrated into release/customer assurance rather than isolated as legal notes.

Three buyer-friction risks:

- An answer may become stale after a provider/model/architecture change.
- Certification language may be overstated if report or certificate scope is not attached.
- AI data-use answers can become ambiguous unless inference, evaluation, product improvement and training are separated.

15.5 Executive pitch

I turn enterprise security questionnaires into a controlled evidence workflow, so the company answers faster without allowing AI or sales teams to invent security claims or reuse stale compliance language.

16 Project 9 - GDPR Article 28 DPA & Subprocessor Governance Engine

16.1 Objective

Operationalize processor-contract requirements and subprocessor governance as an auditable control system.

ICO guidance states that processor contracts must address documented instructions, confidentiality, security, subprocessors, rights assistance, controller assistance, end-of-contract deletion/return and audit/information rights; subprocessors require equivalent downstream protections and authorization. Source: https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/contracts-and-liabilities-between-controllers-and-processors-multi/what-needs-to-be-included-in-the-contract/

17 Portfolio DPA Clause Set

1. Processing Instructions. The Processor shall process Controller Personal Data only on the Controller's documented instructions, including instructions concerning international transfers, unless applicable law requires otherwise. To the extent legally permitted, the Processor shall inform the Controller before processing required by law outside those instructions.

2. Confidentiality. The Processor shall ensure that persons authorized to process Controller Personal Data are subject to appropriate confidentiality obligations and receive access only where required for assigned duties.

3. Security. Taking account of the state of the art, implementation cost, nature, scope, context and purposes of processing, and risk to individuals, the Processor shall maintain appropriate technical and organizational measures designed to protect Controller Personal Data against unauthorized or unlawful processing and accidental loss, destruction, damage, alteration or disclosure.

4. Subprocessors. The Processor shall not appoint a new or replacement Subprocessor without the Controller's prior specific or general written authorization. Under general authorization, the Processor shall provide advance notice identifying the proposed Subprocessor, processing purpose, location and effective date and shall provide a reasonable opportunity to object on legitimate data-protection grounds.

5. Subprocessor Flow-Down. The Processor shall enter into written terms with each Subprocessor imposing data-protection obligations that provide an equivalent level of protection for the obligations applicable to the Processor. The Processor remains responsible to the Controller for the Subprocessor's performance of those obligations.

6. Data Subject Rights. Taking account of the nature of processing, the Processor shall provide appropriate technical and organizational assistance to enable the Controller to respond to applicable requests to exercise data-subject rights.

7. Security and Breach Assistance. Taking account of the nature of processing and information available, the Processor shall assist the Controller with applicable security, breach assessment and notification obligations and shall provide information reasonably necessary for investigation.

8. DPIA and Supervisory Authority Assistance. The Processor shall provide reasonable information and assistance requested by the Controller for data-protection impact assessments and, where applicable, consultation with a competent supervisory authority.

9. Return and Deletion. At the end of the Services, the Processor shall, at the Controller's choice and subject to applicable law, return or delete Controller Personal Data and delete existing copies according to the documented offboarding and backup-deletion schedule.

10. Audit and Information Rights. The Processor shall make information reasonably necessary to demonstrate compliance with applicable processor obligations available to the Controller and shall permit and contribute to proportionate audits or inspections subject to reasonable confidentiality, security, scope and scheduling safeguards.

11. International Transfers. Where Controller Personal Data is transferred to a country or recipient requiring a transfer mechanism under applicable data-protection law, the Processor shall implement the required contractual or other lawful mechanism and maintain supporting transfer documentation.

12. Processing Schedule. The DPA processing schedule shall identify subject matter, duration, nature and purpose, types of personal data, categories of data subjects, and the Controller's rights and obligations. Material changes to processing require an updated written schedule.

These clauses are a proof-of-work drafting sample and require legal review before use in a binding agreement.

17.1 Subprocessor register fields

- Provider legal name
- Service
- Processing purpose
- Personal-data categories
- Data-subject categories
- Hosting/processing locations
- Transfer mechanism
- DPA/SCC status
- Security assurance
- AI/model-provider flag
- Training/retention terms
- Change-notice date
- Customer objection status
- Current risk rating

17.2 Operational control workflow

1. Procurement identifies a vendor that processes personal data.
2. Privacy determines controller/processor role.
3. TPRM assigns risk tier and gathers evidence.
4. Legal confirms Article 28 terms and international-transfer mechanism.
5. Subprocessor is added to the register.
6. Customer authorization/notice process is completed when required.
7. Renewal review validates vendor, purpose, region, assurance and processing changes.
8. Offboarding confirms deletion/return and register closure.

17.3 Executive pitch

> I convert Article 28 from contract text into a controlled operating process: role classification, processor terms, subprocessor approval, evidence, change notification, deletion and audit readiness.

18 Project 10 - Continuous Control Monitoring & Audit Operations System

18.1 Objective

Demonstrate continuous audit-readiness: recurring evidence, test populations, samples, exceptions, remediation and management visibility.

18.2 Evidence collection rules

1. Evidence must cover the stated control period.
2. Screenshots without date/context are insufficient when system export or configuration evidence is available.
3. Evidence must be attributable to a system and owner.
4. Manual evidence must be approved or reviewer-verifiable.
5. Reused evidence must be revalidated for the new period.
6. Sensitive reports are stored under access control and shared externally only through approved channels.
7. Missing evidence creates an exception; it is not automatically evidence that the control failed.

18.3 Audit request lifecycle

Request -> Owner Assignment -> Evidence Collection -> GRC QA -> Auditor/Reviewer Submission -> Follow-Up -> Exception -> Remediation -> Retest -> Closure

18.4 Control test record

Every test record includes:

- test ID;
- control ID;
- test objective;
- period;
- population definition;
- sample selection;
- procedure;
- evidence examined;
- expected result;
- actual result;
- exception/severity;
- remediation owner and due date;
- retest evidence;
- final conclusion.

18.5 Completed 15-request audit simulation

The companion workbook P10_Audit_Ops includes fifteen audit requests spanning logical access, termination, privileged access, SIEM coverage, vulnerabilities, secure change, enterprise risk, TPRM, Article 28, AI governance, Article 50 release evidence, customer assurance, disaster recovery, incident response and encryption.

Two simulated exceptions are intentionally included:

- one production service lacks centralized application logging;

- one EU chatbot Article 50 decision exists but is not linked to release evidence.

This demonstrates exception handling rather than presenting an unrealistically perfect environment.

18.6 Continuous monitoring escalation

A control becomes **At Risk** when:

- evidence is overdue by more than five business days;
- a required control owner is vacant;
- more than 5% of quarterly evidence items are overdue;
- a critical/high exception passes remediation due date;
- a production change materially affects the control without review;
- repeated exceptions indicate a systemic failure rather than an isolated miss.

18.7 Monthly audit-operations dashboard

- Evidence completion percentage
- Overdue evidence items
- Open audit requests
- Auditor follow-ups aging
- Open exceptions by severity
- Overdue remediation
- Controls without current owner
- Control tests completed
- Repeat exceptions
- Customer security questionnaires waiting for evidence

18.8 Executive pitch

I keep GRC audit-ready continuously by defining what evidence is due, who owns it, how it is tested, what counts as an exception and how remediation is retested, instead of treating audit preparation as a once-a-year scramble.

19 LinkedIn Deployment Suite

19.1 Featured section - exact order

1. **Enterprise Trust & Customer Assurance Engine**
   Description: Public-source enterprise procurement evidence analysis showing how security, privacy and AI claims map to buyer questions, evidence and remediation.
   Media: 01_Enterprise_Trust_Customer_Assurance.pdf
   Skills: GRC, Customer Assurance, SOC 2, ISO 27001, AI Governance

2. **AI Governance Operating System - NIST AI RMF + ISO 42001**
   Description: 15-system AI inventory with ownership, risk scoring, human oversight,

evaluation controls and monitoring.
Media: 02_AI_Governance_Operating_System.pdf
Skills: AI Governance, NIST AI RMF, ISO 42001, Risk Assessment

3. **Third-Party Risk & AI Subprocessor Engine**
   Description: Tiered vendor-risk model covering enterprise SaaS, cloud and AI providers, with a 20-question vendor questionnaire.
   Media: 03_TPRM_AI_Subprocessor_Engine.pdf
   Skills: TPRM, Vendor Risk, GDPR, Third-Party Risk

4. **SOC 2 + ISO 27001 Control-to-Evidence System**
   Description: Fifteen-domain control inventory with evidence requirements, frequency, owners and audit test procedures.
   Media: 04_SOC2_ISO27001_Control_Evidence.pdf

5. **EU AI Act Article 50 Transparency Audit System**
   Description: Fifteen-use-case transparency register with disclosure, synthetic-content, deepfake and provenance controls aligned to live 2026 Article 50 obligations.
   Media: 06_AI_Transparency_Article50.pdf

19.2 Skills to add

Governance, Risk Management, and Compliance (GRC); Information Security Governance; Risk Assessment; Third-Party Risk Management; Vendor Risk Management; Technology Risk; Security Compliance; SOC 2; ISO 27001; NIST; AI Governance; Artificial Intelligence Risk Management; Data Privacy; GDPR; Control Assessment; Internal Controls; Risk Management; Security Audits; Policy Development; Regulatory Compliance.

19.3 Ten LinkedIn posts

19.3.1 Post 1 - Enterprise Trust

Most SaaS security friction is not a missing control. It is a missing evidence path. I built an Enterprise Trust & Customer Assurance Engine that maps buyer questions to evidence, framework anchors, owners and remediation. The strongest lesson from the public Lexi case study: two individually reasonable AI-data statements can still create procurement friction when buyers must reconcile retention or training language across multiple documents. The system classifies every conclusion as observed, not publicly identified, or internal verification required - so it improves trust without pretending a website review is an audit. Portfolio artifact: 15-domain evidence matrix + executive remediation roadmap. #GRC #CustomerTrust #SOC2 #ISO27001 #AIGovernance

19.3.2 Post 2 - AI Governance OS

Responsible AI becomes useful when it has owners, evidence and stop conditions. I built a 15-system AI Governance Operating System using NIST AI RMF and ISO 42001. Every system has a purpose, data boundary, human-oversight rule, inherent risk score, controls, residual risk and

monitoring owner. The goal is not a prettier AI policy; it is a repeatable release and risk decision. #AIGovernance #NISTAIRMF #ISO42001 #TechnologyRisk

19.3.3 Post 3 - TPRM

Vendor risk is not a 100-question spreadsheet. It is a decision. I built a TPRM engine covering 10 real enterprise vendors, AI providers and subprocessors. Each record maps data processed, jurisdiction/residency concerns, public security evidence, criticality, required assurance and residual risk. The final state is always Approve, Approve with Conditions, Risk Acceptance or Reject. #TPRM #VendorRisk #GRC #Privacy

19.3.4 Post 4 - Control Evidence

A control framework is not an audit program until someone knows who owns each control, what evidence proves it operated and how it is tested. I built a 15-domain SOC 2 + ISO 27001 control-to-evidence system covering governance, IAM, encryption, logging, vulnerabilities, SDLC, incident response, BCP, vendor risk, privacy and AI governance. #SOC2 #ISO27001 #GRC #AuditReadiness

19.3.5 Post 5 - Executive Risk

Board reporting should answer one question: what decision does leadership need to make? I built an executive technology-risk register with 15 risks, risk appetite, treatments, owners and KRIs. Examples include shadow-AI leakage, stale assurance claims, missed Article 50 controls and model-provider changes without governance review. #RiskManagement #CISO #GRC #BoardReporting

19.3.6 Post 6 - Article 50

EU AI Act Article 50 is live. I built a 15-use-case transparency register that maps interactive AI, synthetic content, deepfakes, public-interest text, emotion recognition and biometric categorization to provider/deployer roles, disclosure controls and provenance evidence. The artifact turns a legal obligation into a release checklist. #EUAIAct #Article50 #AIGovernance #ResponsibleAI

19.3.7 Post 7 - Shadow AI

Shadow AI is a data-loss problem, not just an awareness problem. I built a governance design with an approved-AI allowlist, prompt gateway, classification-aware DLP, local redaction, short-lived memory vault and a SEV-1/2/3 incident playbook for unsanctioned AI egress. #DLP #AIGovernance #InformationSecurity #ShadowAI

19.3.8 Post 8 - Questionnaire Triage

The fastest way to answer security questionnaires badly is to automate unsupported answers. I built a 25-question procurement triage library where every response has required evidence and an accountable owner. AI can draft, but high-risk answers are blocked when evidence is stale or missing. #CustomerAssurance #SecurityQuestionnaire #GRC #EnterpriseSales

19.3.9 Post 9 - Article 28

GDPR Article 28 becomes operational when subprocessors have owners, terms, locations, change notices and deletion evidence. I built a processor/subprocessor engine with 12 clause controls, a vendor register schema and an approval-to-offboarding workflow. #GDPR #Privacy #TPRM #DataProtection

19.3.10 Post 10 - Continuous Audit Ops

Audit readiness is an operating cadence. I built a continuous GRC system that tracks recurring evidence, populations, samples, auditor requests, exceptions, remediation and retest. I intentionally included failed controls in the simulation because a believable audit program shows how exceptions are handled, not a fantasy where every control is green. #Audit #SOC2 #ISO27001 #GRC

20 Visual asset guide - one prompt per project

1. **Enterprise Trust** - “Create a clean executive GRC diagram showing Public Claims -> Buyer Questions -> Evidence -> Control Owner -> Remediation -> Enterprise Trust. White background, navy headings, minimal cybersecurity iconography, no stock-photo people. Overlay: Enterprise Trust Is An Evidence System.”
2. **AI Governance OS** - “Create an enterprise AI governance lifecycle wheel with NIST AI RMF GOVERN at the center and MAP, MEASURE, MANAGE around it; show inventory, owner, risk, evaluation, monitoring. Overlay: Every AI System Needs An Owner And Evidence.”
3. **TPRM** - “Create a vendor-risk funnel from SaaS/AI vendors to risk tiering, evidence review, findings, residual risk and decision. Overlay: Vendor Risk Ends In A Decision.”
4. **Control Evidence** - “Create a control-to-evidence pipeline: Framework -> Control -> Owner -> Evidence -> Test -> Exception -> Remediation. Overlay: Frameworks Do Not Operate Controls. People And Evidence Do.”
5. **Executive Risk** - “Create a board-risk dashboard with five red/amber/green risk cards, arrows showing risk movement and a Decision Required box. Overlay: Risk Reporting Must Drive Decisions.”
6. **Article 50** - “Create an AI transparency decision tree for interactive AI, synthetic content, deepfakes, public-interest text, emotion recognition and biometric categorization. Overlay: Article 50 Is Now A Release-Control Problem.”
7. **Shadow AI** - “Create a data-flow architecture: User -> AI Prompt Gateway -> DLP/Redaction -> Approved AI Provider -> Human Review -> Business App, with blocked path to consumer AI. Overlay: Shadow AI = Data Egress Risk.”
8. **Questionnaire Triage** - “Create a 100-question security questionnaire transforming into canonical answers, evidence links and approvals. Overlay: Automate Drafting, Never Automate Unsupported Claims.”
9. **Article 28** - “Create a controller -> processor -> subprocessor chain with contracts, authorization, equivalent obligations, deletion and audit evidence. Overlay: Article 28 Is A Supply-Chain Control.”

10. **Continuous Audit Ops** - “Create a monthly evidence calendar feeding an audit tracker, exception queue and retest loop. Overlay: Audit Readiness Is A Cadence, Not A Sprint.”

21 Personal Web Portfolio - Page-by-Page Copy

22 Home

22.1 Hero

Technology GRC, Third-Party Risk & AI Governance - demonstrated through operating artifacts.

I build risk registers, control-to-evidence systems, vendor assessments, customer assurance workflows and AI governance artifacts that translate frameworks into decisions a security leader, auditor, buyer or founder can inspect.

Primary CTA: View the 10-project GRC & AI Governance Portfolio
Secondary CTA: Request a 24-Hour Public Trust Readiness Snapshot

22.2 Proof strip

- 10 completed GRC / TPRM / AI governance projects
- 15-system AI governance inventory
- 10-vendor TPRM risk register
- 15-domain SOC 2 / ISO control matrix
- 25-question enterprise assurance knowledge base
- Live Article 50 transparency implementation model

22.3 Integrity statement

Portfolio case studies clearly distinguish public evidence from simulated enterprise controls. I do not present public-source reviews as certifications, penetration tests, legal opinions or client work when no client relationship exists.

23 Portfolio

Use a ten-card grid. Each card contains: business problem, artifact, framework, one key lesson and a link to the full project.

23.1 Card copy

Enterprise Trust & Customer Assurance - Reduce security-review friction by mapping buyer questions to current evidence and accountable owners.

AI Governance Operating System - Govern AI systems through inventory, risk, human oversight, testing, monitoring and change control.

TPRM & AI Subprocessors - Turn vendor onboarding into a risk-tiered decision with evidence and residual-risk treatment.

SOC 2 + ISO Control Evidence - Translate frameworks into control statements, owners, evidence and test procedures.

Executive Technology Risk - Convert compliance activity into risk appetite, KRIs and decisions for leadership.

AI Transparency & Article 50 - Turn live EU AI transparency duties into product/release controls.

Shadow AI & Prompt DLP - Protect confidential data with approved AI channels, local redaction and incident response.

Procurement Security Triage - Answer buyer questionnaires faster without reusing stale or unsupported claims.

GDPR Article 28 Engine - Operationalize processor contracts and subprocessor governance.

Continuous Audit Operations - Maintain evidence, exceptions, remediation and retest throughout the year.

24 Services

24.1 24-Hour AI & Security Public Trust Readiness Snapshot - $99 founding-client scope

A source-backed review of public security, privacy and AI trust evidence for B2B SaaS/AI teams preparing for enterprise procurement.

Includes:

- 15-point public evidence matrix
- buyer-question register
- third-party/AI risk observations
- five priority documentation actions
- executive PDF

Scope is public evidence only. It is not a certification, penetration test, legal opinion or formal compliance audit.

25 About

I focus on the operating layer of Technology GRC and AI governance: what the risk is, which control addresses it, what evidence proves operation, who owns it, what remains unresolved and what decision the business needs to make.

My portfolio is built to be inspectable. Every major artifact includes control ownership, evidence, testing logic, risk treatment or implementation status rather than framework-name decoration.

26 Contact

Best fit: B2B SaaS/AI founders, GRC/security teams, third-party risk teams, customer assurance teams and remote GRC/technology-risk hiring managers.

Call to action: Send the word **TRUST** for the fixed-scope Public Trust Readiness Snapshot, or send **GRC** for the portfolio index.

27 Global Contractor Positioning & Compensation Scripts

27.1 Critical boundary

A contractor arrangement does not automatically eliminate immigration, tax, sanctions, export-control, worker-classification, data-access, residency or employer policy constraints. Never claim that B2B contracting legally bypasses a restriction. The correct positioning is that a defined independent service can sometimes be purchased cross-border when the buyer's legal/procurement policies permit it.

27.2 B2B contractor framing

I work on defined Technology GRC, TPRM and AI-governance deliverables as an independent remote service provider. The engagement is scoped around outputs, evidence, deadlines and acceptance criteria rather than employee headcount. I can work asynchronously across time zones and deliver through the client's approved collaboration and security processes. Before engagement, I am happy to complete the client's vendor onboarding, confidentiality, data-protection and security requirements. If the company's legal or procurement policy restricts cross-border contractors, I respect that restriction rather than asking the company to work around it.

27.3 Experience objection

I do not represent the portfolio as years of employment I have not held. What I can show directly is the work product: AI system inventories, risk registers, TPRM assessments, control-to-evidence mapping, questionnaire governance, Article 50 transparency controls and continuous audit operations. I would prefer that you evaluate those artifacts and my reasoning against the scope you need. For responsibilities that require historical ownership of a live certification audit, regulator interaction or production security program, I would not claim that a portfolio substitutes for that operating history.

27.4 Certification objection

I treat certifications as signals, not substitutes for control operation. My current evidence is the completed workflow itself: map the requirement, design the control, name the owner, identify the evidence, test the result, record exceptions and drive remediation. If a specific credential is genuinely required for the engagement, I will treat that as an eligibility requirement rather than claim the portfolio replaces it.

27.5 Rate positioning

Current 2026 U.S. job postings show wide senior GRC compensation ranges, including roughly $100,000-$185,000+ base depending on scope and market, while one remote 1099 GRC/ISSO posting listed $55-$62.50/hour depending on experience. High independent rates such as $75-$150/hour are therefore a later commercial target for specialized, outcome-based work - not a defensible default rate solely because these portfolio artifacts exist.

27.5.1 Response when asked for hourly rate

For a defined deliverable I prefer fixed scope because it aligns cost with the outcome. For advisory or open-ended support, my rate depends on the responsibility, data sensitivity, turnaround, meeting load and whether I am only producing analysis or owning implementation. I price after scope. I do not anchor a senior consulting rate solely on a title; the rate needs to match the value and accountability of the work.

27.5.2 Response if pressed for a range

For narrowly scoped analyst execution, I would price below senior specialist consulting rates while I build additional live client ownership. For specialized, time-sensitive work where the client is buying a finished TPRM assessment, customer-assurance package or AI-governance implementation rather than training me, I would quote the project based on effort and business value. I would only move toward $75-$150/hour equivalents after real paid engagements and repeatable demand support that level.

27.6 Founder/CTO cold pitch

I reviewed your public security, privacy and AI trust footprint from an enterprise procurement perspective. The issue I look for is not “are you compliant?” - it is where a buyer still has to email engineering or legal because the public evidence is unclear, inconsistent or difficult to find.

I offer a fixed-scope 24-Hour AI & Security Public Trust Readiness Snapshot. It maps 15 buyer-facing evidence areas, identifies the highest-friction unanswered questions, distinguishes public evidence from internal verification, and gives five concrete documentation actions. The founding-client price is $99 USD upfront; it is not a certification, penetration test or legal opinion.

If useful, reply **YES** and I will send the secure payment link. The report is delivered after confirmed payment, with one factual correction/revision included.

27.7 Salary conversation for employment

I would like to separate level from aspiration. I am not asking to be levelled as a senior employee solely because I built senior-style artifacts. I am asking to be evaluated against the actual responsibilities and quality bar of the role. If the role is analyst-level, I would expect compensation consistent with the company's market band. If the scope includes independent ownership of audits, enterprise risk programs, customer assurance and AI governance, then I would expect the level and compensation to reflect that broader accountability once I can demonstrate it in production.

28 Source Register

All web sources are official or primary sources where available.

- **NIST AI RMF:** https://www.nist.gov/itl/ai-risk-management-framework
- **NIST GenAI Profile:** https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
- **NIST AIRC Core:** https://airc.nist.gov/airmf-resources/airmf/5-sec-core/
- **EU Article 50 Guidelines:** https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems
- **EU Transparency Code:** https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content
- **ISO 27001:** https://www.iso.org/standard/27001
- **ISO 42001:** https://www.iso.org/standard/42001
- **AICPA TSC:** https://www.aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022
- **ICO Article 28:** https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/contracts-and-liabilities-between-controllers-and-processors-multi/what-needs-to-be-included-in-the-contract/
- **OpenAI Enterprise Privacy:** https://openai.com/enterprise-privacy/
- **AWS Bedrock Security:** https://aws.amazon.com/bedrock/security-privacy-responsible-ai/
- **Gemini Enterprise Security:** https://docs.cloud.google.com/gemini/enterprise/docs/compliance-security-controls
- **Slack Security:** https://slack.com/trust/security
- **GitHub Data Residency:** https://docs.github.com/en/enterprise-cloud@latest/admin/data-residency/about-storage-of-your-data-with-data-residency
- **Lexi Terms:** https://www.getlexi.io/terms
- **Lexi Data Controls:** https://www.getlexi.io/data-controls
- **Lexi Privacy:** https://www.getlexi.io/privacy
- **Lexi Home:** https://www.getlexi.io/

28.1 Compensation evidence used in positioning

- Modine Senior GRC Analyst, posted July 2026, estimated $100,000-$140,000, U.S.-remote with U.S. work authorization requirement: https://careers.modine.com/job/Senior-GRC-Analyst/9097-en_US/
- Astra GRC Analyst, U.S.-remote, $95,000-$135,000: https://jobs.ashbyhq.com/astra/07b241b9-b5c2-49cd-9687-60c7c21707f1
- Compyl Senior GRC Analyst, U.S.-remote, $115,000-$145,000: https://jobs.ashbyhq.com/compyl/62d91d65-099c-489c-baab-d5f0be149eca
- Doppler Senior GRC Analyst, U.S.-remote, $150,000-$185,000: https://jobs.ashbyhq.com/doppler/7dcb6231-86c0-4f84-b9fb-0f2ffd1d3117
- SecureSoft GRC Analyst / ISSO listing included a 1099 hourly range of $55-$62.50 depending on experience and U.S.-specific eligibility requirements: https://www.indeed.com/viewjob?jk=f97b0b38de3359ec