PRACTITIONER TOOLKIT · 2026

AI Compliance Toolkit

25 Templates & Checklists

Toolkit · Templates · Checklists · Skills

DPIA | RISK REGISTER | AUDIT CHECKLISTS | NIST AI RMF | EU AI ACT | ISO 42001 | MODEL RISK | GENAI FRAUD

Inside this toolkit

- 25 production-ready templates — DPIA, framework, risk register, audits
- Frameworks: NIST AI RMF · EU AI Act · ISO 42001
- Model risk management: lineage, drift, validation, attestation
- Generative AI fraud detection workflow blueprints (banking)

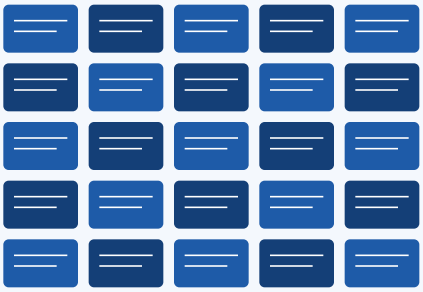

25 Templates Inside
DOCX · XLSX · PDF · ready to deploy

INSIDE THE TOOLKIT

What you'll find in 20 pages

This toolkit consolidates 25 production-ready templates, checklists, and workflow blueprints that AI Compliance Officers, Model Risk teams, and Internal Audit functions can deploy immediately. Every artifact is mapped to at least one external framework: NIST AI RMF 1.0, the EU AI Act, or ISO/IEC 42001.

25
TEMPLATES & CHECKLISTS
DOCX, XLSX, PDF formats

3
FRAMEWORKS MAPPED
NIST · EU AI Act · ISO 42001

12+
WORKFLOW BLUEPRINTS
Banking, healthcare, public sector

Table of contents

ORIENTATION

How to use this toolkit

Use this toolkit as a starting library, not a finished product. Each template is intentionally generic; adapt the language to your sector, regulator, and risk appetite before adopting. We recommend running through the 90-day implementation roadmap on page 19 the first time you stand up an AI governance program.

Operating model — three stages

01 DISCOVER

Inventory AI use cases. Classify by risk tier.

USE:

- Template 03 — AI Use Case Inventory
- Template 17 — Risk Tiering Worksheet

02 ASSESS

Run DPIAs, AIA, bias & explainability tests.

USE:

- Template 05 — DPIA core form
- Template 21 — Bias & fairness test plan

03 OPERATE

Monitor in production. Attest. Respond to incidents.

USE:

- Template 11 — Drift detection log
- Template 24 — Incident response runbook

A note on adaptation

Templates are deliberately conservative. If your organisation operates under sector-specific rules (e.g. SR 11-7 for US banks, MAS FEAT for Singapore financial institutions, or NHS DSPT for UK healthcare), layer those obligations on top — do not replace this baseline.

TEMPLATES 1–4

AI compliance foundations

The first four templates establish the spine of an AI compliance program. Deploy these before anything else — every later artifact references them.

01 AI Compliance Framework

Master document defining scope, policy hierarchy, governance bodies, and decision rights for AI.

02 AI Risk Register

Live register tracking inherent risk, residual risk, controls, and owner per AI system.

03 AI Use Case Inventory

Searchable catalog of all AI systems in scope with metadata: purpose, data, risk tier.

04 AI Governance Board Charter

Charter, membership, quorum rules, and escalation pathways for the AI governance body.

Why these four come first

Auditors and regulators look for evidence of structure before evidence of activity. A board charter and a live risk register answer 80% of any 'show us your AI governance' opening request — even before specific controls are deployed.

TEMPLATES 5–8

DPIA & impact assessments

DPIAs (GDPR Article 35) and Algorithmic Impact Assessments remain the most-requested deliverables across audits in 2026. These four templates are designed to satisfy both UK ICO and EU EDPB guidance without duplication.

05 DPIA — core form
Standard data protection impact assessment template aligned with GDPR Art. 35 and ICO guidance.

06 AIA — Algorithmic Impact Assessment
Canada Treasury Board-style AIA covering automation level, affected populations, redress.

07 DPIA threshold trigger checklist
Twelve-question screening tool that determines whether a full DPIA is required.

08 Data minimisation worksheet
Per-field justification grid mapping each personal data field to its lawful basis.

DPIA section structure — at a glance

Section	Coverage
§1 Context	Purpose · scope · data flows · stakeholders
§2 Necessity	Lawful basis · proportionality · alternatives considered
§3 Risks to data subjects	Identified harms · likelihood · severity rating
§4 Controls	Technical · organisational · residual risk after controls
§5 Consultation	DPO opinion · regulator consultation if required
§6 Sign-off	Approver · review cadence · trigger conditions for re-assessment

TEMPLATES 9–12

Audit & assurance checklists

Used by internal audit, second-line risk teams, and external assurance providers. Each checklist is constructed so a non-technical auditor can use it without ML literacy.

09 Pre-deployment audit checklist

32-item gate review covering documentation, testing, sign-offs, and rollback plans.

10 Production audit checklist

Quarterly review covering drift, performance, incident history, and re-training events.

11 Drift detection log

Template log capturing input drift, output drift, and concept drift signals per model.

12 Validation report template

Independent model validation report — methodology, findings, conditions, conclusions.

TEMPLATES 13–16

Model risk management

Banks and insurance carriers built mature model risk management (MRM) practices well before AI emerged. These templates extend established MRM hygiene to ML and generative AI systems.

13 Model lineage tracker

End-to-end trace: training data → preprocessing → model artifact → deployment → version.

14 Drift monitoring plan

Monitoring strategy: metrics, thresholds, alert routing, retraining triggers.

15 Senior management attestation

Quarterly attestation letter — model owner confirms continued fitness-for-purpose.

16 Model card template

Standardized model card capturing intended use, training data, performance, limitations.

Three-lines model for AI — responsibilities

Line	Who	What they do
1st line	Model developers, data science teams	Build & document; self-test before submission
2nd line	Model risk / AI compliance	Independently validate; challenge assumptions
3rd line	Internal audit	Test the validation; report to audit committee

TEMPLATES 17–20

Vendor & third-party AI

Most enterprises consume more AI than they build. These four templates harden the procurement and ongoing-monitoring of third-party AI — including foundation-model APIs.

17 Vendor AI due diligence questionnaire

76-question RFP supplement covering data, training, IP, security, and exit terms.

18 Contractual AI clauses library

Negotiated clauses: warranties, audit rights, model change notice, indemnities.

19 Foundation model API risk assessment

Specific to LLM-as-a-service: prompt injection, output filters, residency.

20 Sub-processor & supply-chain map

Visual map identifying every AI sub-processor and where data crosses borders.

FRAMEWORKS AT A GLANCE

AI risk assessment comparison matrix

The three frameworks below are the most commonly referenced across audits in 2026. Most organisations need to map their controls to two or three simultaneously.

Dimension	NIST AI RMF 1.0	EU AI Act	ISO/IEC 42001
Origin	NIST (United States)	European Union	ISO/IEC (international)
Status	Voluntary guidance	Binding regulation	Voluntary standard
Structure	4 functions: Govern, Map, Measure, Manage	Risk-tiered: prohibited / high / limited / minimal	AI Management System (annex A controls)
Best for	US federal contractors, voluntary alignment	Anyone selling AI into the EU	Multinationals seeking certifiable AIMS
Certifiable?	No (self-attest)	Yes (CE-mark for high-risk)	Yes (3rd-party)
Sanctions	Reputational only	Up to 7% global turnover	Loss of certification
Mapping template	T22 in this kit	T23 in this kit	T25 in this kit

70%+
ENTERPRISES MAPPING TO 2+ FRAMEWORKS
Survey of 850+ practitioners, 2026

T22–T25
TEMPLATES THAT HANDLE THE MAPPING
Pre-mapped rows for each framework

NIST AI RMF

Function mapping checklist

NIST AI RMF 1.0 organises practice into four functions. Use this checklist as the first cut when scoping a NIST-aligned program — each item is also referenced from Template 22.

GOVERN

- ✓ Policies, processes, accountability defined
- ✓ Board-level oversight documented
- ✓ Roles & responsibilities mapped

MAP

- ✓ Context established (purpose, stakeholders)
- ✓ Categorisation by impact tier
- ✓ Risks identified across the lifecycle

MEASURE

- ✓ Quantitative & qualitative metrics defined
- ✓ Independent testing & evaluation in place
- ✓ Trustworthy AI characteristics assessed

MANAGE

- ✓ Risks prioritised and treatment chosen
- ✓ Incident response in production
- ✓ Continuous improvement loop closed

EU AI ACT

Risk-tier compliance map

The EU AI Act applies a risk-tiered structure. Use the map below to determine which obligations apply to a given system — Template 23 contains the full clause-by-clause mapping.

Risk tier	Typical examples	Obligation summary
Prohibited	Social scoring; real-time biometric ID in public spaces (with exceptions)	Banned outright
High-risk	Hiring, credit scoring, medical devices, critical infrastructure, education	Full conformity assessment, CE-mark, post-market monitoring
Limited-risk	Chatbots, deepfakes, emotion recognition	Transparency obligations only
Minimal-risk	Spam filters, AI in video games, inventory optimization	Voluntary codes of conduct
General-purpose AI	Foundation models above compute thresholds	Documentation, copyright policy, evaluation reporting

What 'high-risk' actually triggers

If a system lands in the high-risk tier, expect to deliver: (1) a risk management system, (2) data governance documentation, (3) technical documentation per Annex IV, (4) record-keeping, (5) transparency to deployers, (6) human oversight design, (7) accuracy / robustness / cybersecurity evidence, and (8) a quality management system.

ISO/IEC 42001

Control library — by clause

ISO/IEC 42001 specifies an AI Management System (AIMS), modelled structurally on ISO 27001. The clause-level summary below anchors Template 25 — the auditable control library.

§	Clause	Coverage
4	Context of the organisation	Internal/external issues; interested parties; scope
5	Leadership	Policy; roles; commitment from top management
6	Planning	Risks & opportunities; AI objectives; planning of changes
7	Support	Resources; competence; awareness; documented information
8	Operation	Operational planning; AI risk assessment; AI impact assessment
9	Performance evaluation	Monitoring; internal audit; management review
10	Improvement	Nonconformity & corrective action; continual improvement
A	Annex A controls	AI policies; data; lifecycle; impact; reuse; communication

GEN-AI FRAUD DETECTION

Architecture blueprint (banking)

Generative AI is now a first-line tool for transaction monitoring, synthetic identity detection, and adversarial fraud-pattern discovery. The architecture below is the banking-grade reference used in Templates 21 and 22.

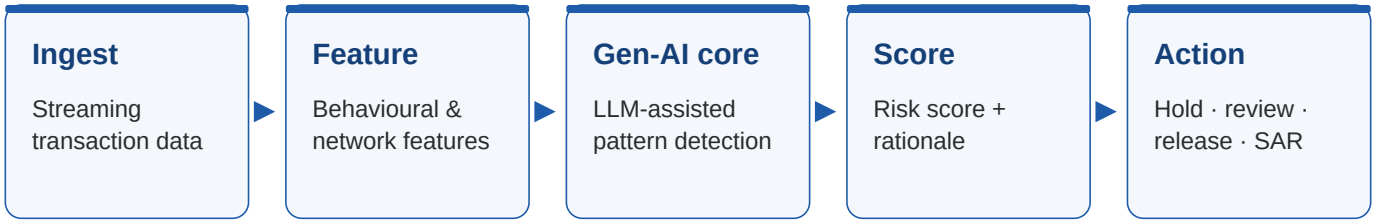

Controls embedded at each stage

Stage	Compliance controls
Ingest	Lineage capture (T13); PII tagging at ingest
Feature	Feature-store versioning; bias monitoring on protected attributes
Gen-AI core	Prompt audit log; output filters; rate-limited LLM calls
Score	Score explanation captured (SHAP / rationale); human-in-loop for borderline
Action	Two-person approval for high-value holds; SAR template T24

GEN-AI FRAUD DETECTION

Workflow & sample event

Below is one end-to-end event traced through the architecture. Each row in the workflow has a corresponding entry in the audit log template (T11). Use this as your training scenario when onboarding new compliance hires.

Timestamp	Stage	Event
T+00:00:00	Ingest	Wire transfer initiated: $94,210, beneficiary in new corridor
T+00:00:01	Feature	12 behavioural features computed; velocity anomaly flagged
T+00:00:02	Gen-AI core	LLM matches pattern to known mule-network typology #47
T+00:00:03	Score	Risk score: 0.87 (high); rationale: corridor + velocity + entity
T+00:00:04	Action	Auto-hold; routed to L2 investigator queue
T+00:18:23	Review	L2 reviewer confirms typology; escalates to MLRO
T+01:42:11	Decision	MLRO approves SAR filing; customer informed within reg window
T+24h	Feedback	Outcome label fed back to feature store; model card updated

BANKING-GRADE CONTROLS

Supplementary requirements

Banking deployments inherit additional obligations from supervisory expectations. The supplement below maps the toolkit templates to the most-cited supervisory frameworks.

Supervisory framework	Focus area	Toolkit map
US — SR 11-7 (Federal Reserve)	Model risk management	T12 · T13 · T14 · T15
EU — EBA/GL/2017/16 (ICAAP)	Internal capital adequacy of AI models	T02 · T11 · T22
UK — PRA SS1/23	Model risk for banks	T12 · T13 · T15 · T25
Singapore — MAS FEAT	Fairness, ethics, accountability, transparency	T21 · T22 · T23
Hong Kong — HKMA AI principles	Governance + customer protection	T01 · T04 · T17
Basel III — operational risk	Loss event reporting for AI failures	T24 — incident response runbook

Common audit finding to avoid

Banking auditors increasingly flag the absence of senior management attestation as a critical finding — not just model documentation gaps. Use Template 15 quarterly even if no regulatory deadline forces it; the audit trail alone usually closes the finding.

TEMPLATES 21–25

Bias, fairness, incident response

The final five templates close out the toolkit and cover the most regulator-scrutinised areas of AI compliance in 2026.

21 Bias & fairness test plan

Protected-attribute test design, fairness metrics, acceptance thresholds, sign-off.

22 NIST AI RMF mapping spreadsheet

Pre-populated rows: each NIST sub-category mapped to internal control IDs.

23 EU AI Act mapping spreadsheet

Clause-by-clause obligations for high-risk + GPAI providers and deployers.

24 Incident response runbook

Detection · triage · containment · notification · post-incident review.

25 ISO/IEC 42001 audit checklist

Clause-by-clause and Annex-A auditor-style review checklist.

MODEL CARD & EXPLAINABILITY

Template structure

Template 16 (model card) and Template 21 (bias test plan) frequently merge into a single deliverable. The structure below is the recommended outline — already field-tested at Big-4 audits.

§	Section	What goes here
1	Model details	Owner, version, framework, training date, lineage link
2	Intended use	Primary uses, out-of-scope uses, intended users
3	Training data	Sources, licensing, consent basis, geographic coverage
4	Performance	Accuracy, precision, recall by sub-group; confidence intervals
5	Fairness analysis	Protected-attribute metrics, thresholds, mitigations applied
6	Explainability	Method (SHAP, LIME, etc.), illustrative examples, limits
7	Caveats	Known failure modes, prohibited use cases, recommended overrides
8	Sign-off	Model owner, validator, AI Compliance Officer, date

Length matters less than auditability

Aim for a model card that fits on three sides of A4 — long enough to satisfy reviewers, short enough that engineers actually maintain it. Anything beyond five pages tends to stop being updated within two release cycles.

INCIDENT RESPONSE

Escalation playbook

Template 24 covers the operational runbook. Below is the escalation ladder organisations should publish alongside it — every AI compliance hire should be able to recite this from memory.

Severity	Trigger criteria	Escalation
Severity 1	Live customer harm, regulator notification trigger, or material financial loss	MLRO + CCO immediately; 24-hour board notification
Severity 2	Material model degradation; possible bias incident; partial customer impact	AI Compliance Officer + Model Owner within 4 hours
Severity 3	Internal-only anomaly; no customer impact; precautionary action needed	Next business day; logged in T24 runbook
Severity 4	Routine; minor drift; covered by automated controls	Captured in weekly monitoring report

IMPLEMENTATION

90-day roadmap

Use the schedule below to stand up an AI compliance baseline in three months. Every milestone references the relevant toolkit templates so a single owner can execute the plan.

Window	Milestone	Templates
Days 1–10	Confirm scope; appoint AI Compliance Officer; stand up governance board	T01 · T04
Days 11–20	Build initial AI use case inventory; categorise by risk tier	T03 · T17
Days 21–35	Run DPIA / AIA on first wave of high-risk systems	T05 · T06 · T07
Days 36–50	Deploy model risk management baseline: lineage, drift, validation	T13 · T14 · T12
Days 51–65	Stand up vendor & 3rd-party AI assessment process	T17 · T18 · T19
Days 66–80	Activate monitoring & incident response capability	T11 · T24
Days 81–90	First quarterly attestation; framework mapping review	T15 · T22 · T23 · T25

What to deliberately defer beyond day 90

Resist the urge to formalise red-teaming, advanced explainability tooling, or jurisdiction-specific filings in the first 90 days. Those are 180-day items — premature deployment of them is the most common reason early AI compliance programs collapse.

HOW TO DEPLOY

Next steps & about GSDC

Each template ships as an editable file (DOCX, XLSX, or PDF). Adapt header, scope statement, and ownership fields to your organisation before adoption. Re-publish under your own governance policy as part of your AI management system.

Deployment checklist

- Confirm executive sponsorship and AI Compliance Officer accountability.
- Adopt the framework mapping that matches your regulatory exposure.
- Customize every template's branding, ownership, and signature blocks.
- Schedule the first attestation cycle within 90 days of adoption.
- Train AI compliance and audit staff on the toolkit's structure.

About GSDC

The Global Skill Development Council is an independent certification body credentialing AI governance, compliance, and risk practitioners. The GSDC AI Compliance Officer Certification and the underlying toolkit are recognized in 180+ countries and used by practitioners at global banks, Big-4 consultancies, FAANG-scale technology firms, and public-sector regulators.