GRC PRACTITIONER REFERENCE · AI GOVERNANCE

NIST AI RMF: Implementation Checklist

GOVERN → MAP → MEASURE → MANAGE: the 4 core functions of the NIST AI Risk Management Framework, mapped to ISO/IEC 42001 Clauses 4–10 and EU AI Act Articles 9–15, 72–73.

→ If you build, review, or audit AI systems, run through each column before your next deployment or governance review.

12–14%	€35M / 7%	12–18 mo	Aug 2025	Aug 2026
OF ORGS HAVE MATURE AI GOVERNANCE IN PLACE	EU AI ACT MAX FINE (GLOBAL REVENUE)	ISO 42001 INITIAL CERTIFICATION TIMELINE	EU AI ACT GPAI OBLIGATIONS ALREADY IN FORCE	ANNEX III HIGH-RISK AI FULL COMPLIANCE DEADLINE

CROSS-REFERENCED WITH NIST AI RMF (2023) ↔ ISO 42001 Cl. 4–10 ↔ EU AI Act Art. 9–15 · 72–73

HOW TO USE

Governance review: Work left to right. GOVERN sets your policy foundation, MAP inventories risks per system, MEASURE defines your evidence, MANAGE closes the loop. Each cycle updates the next.

Compliance mapping: ISO tags (`Cl.X`) = auditable ISO 42001 clauses. EU tags (`Art.X`) = EU AI Act articles. Check both columns before any conformity assessment.

GOVERN

Organizational Culture, Policies & Oversight

Foundation function. Sets the tone, structure, and accountability before any AI system is built or deployed.

- [ ] **Establish an AI governance committee** with defined roles, responsibilities, and escalation paths `Cl.5`
- [ ] **Create an AI policy** that sets boundaries of acceptable use and is signed off at board/executive level `Cl.5`
- [ ] **Document your risk tolerance statement**: what risk appetite applies to AI decisions, automated outputs, and high-risk use cases `Cl.6`
- [ ] **Define AI incident ownership**: who is notified, who decides, who communicates externally when an AI system malfunctions or causes harm `Cl.10`
- [ ] **Designate an AI officer or named responsible person** per high-risk system. Identifiable human accountability is required; governance must not be anonymous `Art.9`
- [ ] **Define governance requirements for third-party AI tools and APIs**. Vendor or open-source AI integrated into your stack is still your liability under EU AI Act `Art.9`

MAPS TO
ISO Cl.4 Context · ISO Cl.5 Leadership · ISO Cl.6 Planning · ISO Cl.10 Improv · Art.9 Risk Mgmt

MAP

AI Use-Case Inventory & Risk Identification

Context function. Identifies what AI systems exist, who they affect, and what could go wrong, per system.

- [ ] **Build an AI system inventory**: all deployed and in-development AI systems, their intended purpose, and deployment context `Cl.4`
- [ ] **Identify stakeholders per system**: users, affected parties (including non-users), and applicable regulators `Cl.4` `Art.9`
- [ ] **Map data sources and lineage**: where does training/inference data come from, how was it collected, what pre-processing was applied `Art.10`
- [ ] **Perform impact assessments** for high-risk systems. Document the negative impacts if the system fails or behaves unexpectedly `Cl.8` `Art.9`
- [ ] **Create risk registers** and threat models per AI system. Document known and foreseeable risks to health, safety, and fundamental rights `Cl.6`
- [ ] **Classify systems by EU AI Act risk tier** (Unacceptable / High-Risk / Limited / Minimal). Annex III defines what counts as high-risk `Art.6 / Annex III`

MAPS TO
ISO Cl.4 Context · ISO Cl.6 Planning · ISO Cl.8 Operation · Art.6 / Annex III · Art.9 Risk Mgmt · Art.10 Data Gov

MEASURE

Testing, Metrics & Monitoring Evidence

Validation function. Defines measurable targets and generates the evidence that the system works as intended.

- [ ] **Define accuracy metrics per system**: precision, recall, F1, AUROC. Declare them in technical documentation, not just internally `Art.15`
- [ ] **Run bias evaluations across demographic groups**: an 80% accurate hiring AI might be 90% for one group and 70% for another; document the disparity `Art.10`
- [ ] **Conduct adversarial robustness testing**: generate adversarial examples, test for jailbreaks, measure out-of-distribution input handling `Art.15`
- [ ] **Implement production monitoring** for model drift as part of your formal post-market monitoring program. Data distributions change; accuracy from development does not hold indefinitely `Art.9` `Art.72`
- [ ] **Publish model cards and data sheets** with testing results, validation scenarios, and accuracy trade-offs `Cl.9` `Art.11`
- [ ] **Schedule internal audits and management reviews**. Collect evidence that the AI management system is working effectively `Cl.9`
- [ ] **Verify AI outputs are interpretable by end users**. Document instructions for use; high-risk AI must enable operators to correctly interpret and act on outputs, not just produce accurate ones `Art.13`

MAPS TO
ISO Cl.9 Eval · Art.9 Risk Mgmt · Art.10 Data Gov · Art.11 Tech Docs · Art.13 Transparency · Art.15 Accuracy · Art.72 Post-Market

MANAGE

Incident Response & Continuous Improvement

Response function. Implements controls, handles incidents, and feeds learnings back into GOVERN and MAP.

- [ ] **Implement risk treatment controls**: design controls, training requirements, and operational procedures for identified risks `Cl.8` `Art.9`
- [ ] **Deploy human oversight mechanisms**: human-in-the-loop (HITL), human-on-the-loop (HOTL), or human-in-command (HIC). A kill switch is required for high-risk systems `Art.14`
- [ ] **Establish immutable audit logging**: log operating period, input data, query databases, verifying persons. Retain logs with integrity controls; tamper-evident storage or cryptographic signing is recommended implementation practice `Art.12`
- [ ] **Define vulnerability patching SLAs**: industry baseline is 14 days for critical CVEs, 30 days for high-severity. Track ML framework dependencies separately from your standard patch pipeline `Art.15`
- [ ] **Conduct AI-specific red teaming**: attempt jailbreaks, training data extraction, input poisoning. This is distinct from traditional pen testing `Art.15`
- [ ] **Run corrective action on non-conformities**: every incident and audit finding feeds back into updated GOVERN policies and MAP threat models `Cl.10`
- [ ] **Report serious AI incidents to national market surveillance authority**. Art.73 deadlines are tiered: 2 days for critical infrastructure disruptions, 10 days if a death is involved, 15 days for all other serious incidents. Define severity thresholds and notification owners now `Art.73`

MAPS TO
ISO Cl.8 Ops · ISO Cl.10 Improv · Art.9 Risk Mgmt · Art.12 Logging · Art.14 Oversight · Art.15 Security · Art.73 Incidents

⚠ WHAT MOST ORGS SKIP

Only 12–14% of organizations have enterprise-level AI governance in place. Here's why:

× Governance as a one-time event
Teams run a risk assessment at launch and never touch it again. The NIST AI RMF is a continuous cycle: each MAP → MEASURE → MANAGE pass must feed back into updated GOVERN policies. That feedback loop is where most programs fall apart.

× No per-system AI inventory
Most orgs treat AI governance as a single category policy, not per deployment. ISO 42001 Clause 4 and EU AI Act Article 11 require per-system documentation (intended purpose, affected populations, deployment context) before any conformity assessment.

× Accuracy without fairness checks
A model with 80% overall accuracy can be 70% accurate for specific demographic groups. EU AI Act Article 10 requires examining datasets for statistical biases. Poor data quality is the most common compliance gap.

KEY INSIGHTS

INSIGHT 01
GOVERN → MAP → MEASURE → MANAGE is a loop, not a checklist
Each cycle ends by feeding audit findings and incident learnings back into updated governance policies and new risk mappings. Skipping this feedback loop is the single most common governance failure.

INSIGHT 02
ISO 42001 Clauses 4–10 align directly to all four NIST functions
If you're ISO 27001-certified, the Annex SL structure is identical. Clauses 5–6 map to GOVERN, 8 maps to MAP and MANAGE, and 9–10 maps to MEASURE and Improvement. Integration saves significant effort.

INSIGHT 03
EU AI Act Article 15 requires AI-specific security, not just standard appsec
Input validation, output filtering, adversarial robustness, and prompt injection defenses are Article 15 requirements. Traditional pen testing does not cover jailbreaking or training data extraction. You need a separate AI red team exercise.