

■ MercurySecurity — 30-Day AI Governance Sprint Checklist

This checklist breaks down the 4-week sprint we run with clients to establish a credible, evidence-based AI governance foundation. Use it as a self-assessment or to prepare for a sprint with us.

Week 1 — Scope & Map

- Identify high-risk AI systems (those impacting hiring, lending, healthcare, or critical services).
- Define intended purpose & out-of-scope use.
- Map system architecture and dependencies.
- Document data sources and provenance.
- Identify responsible owners and stakeholders.

Week 2 — Evidence & Baselines

- Gather training & validation datasets; record lawful basis (GDPR, contracts, consent).
- Run fairness, bias, robustness, and accuracy tests.
- Capture current model performance metrics.
- Document existing oversight processes.
- Establish a secure, version-controlled Evidence Repository.

Week 3 — Controls & Oversight

- Draft a Responsible AI Policy (roles, escalation, prohibited uses).
- Define human-in-the-loop checkpoints and decision boundaries.
- Create a Model Change Log (training, retrains, deployments).
- Draft an Incident & Escalation Playbook.
- Align with EU AI Act, NIST AI RMF, GDPR, ISO/IEC 42001 controls.

Week 4 — Roadmap & Communication

- Assemble Evidence Pack v1.0 (data lineage, tests, policies, logs).
- Deliver a board-ready 90-day governance roadmap.
- Prepare an executive briefing deck for board & investors.
- Set up monthly monitoring cadence.
- Plan for quarterly Evidence Pack refresh.

Next Steps

- Download the Case Study: <https://yourdomain.com/case-study-download/>
- Book a 20-min Governance Briefing: <https://yourdomain.com/?fluent-booking=calendar&host;=mercer&event;=45min>