The Metrics That Actually Matter

Forget vanity metrics. The KPIs that show control (and progress) look like this:

***%** of AI apps/services discovered & risk-ranked*

***%** of AI workflows with guardrails/testing*

***#** of AI data leakage / policy-violation incidents*

(track trends; add time-to-detect/ time-to-contain if you want maturity proof)

If those numbers are unknown, the program hasn't started yet.

ACTION BOX:

The First 90 Days (No Mythical "AI Strategy" Required)

DAYS 1–30:

BUILD AN AI ASSET INVENTORY

Inventory everything AI touches:

- Models (hosted/internal/fine-tuned)
- Copilots and AI assistants
- Plugins/tools/agents and their permissions
- Datasets and knowledge bases
- AI-enabled workflows across departments
- Systems these workflows can read/write (email, CRM, cloud storage, code repos)

Deliverable:
A living AI inventory with owners + risk ranking.

DAYS 31–60:

PUT CONTROLS WHERE THE MONEY IS

Don't boil the ocean. Start with the **top three highest-risk use cases** (sensitive data + high adoption + high impact).

Add guardrails:

- Least-privilege tool access for agents
- Human approval for sensitive actions
- Data boundaries (what can be prompted, stored, retrieved)
- Logging tied to identity + workflow context

Deliverable:
Three "lighthouse" workflows that are measurably safer.

DAYS 61–90:

AUTOMATE AI SECURITY TESTING + VENDOR ONBOARDING

Make AI security repeatable:

- Automated testing in pipelines for injection and leakage behaviors
- Policy checks for prompts/outputs/logging
- Vendor onboarding requirements for AI features (data retention, training use, auditability)

Deliverable:
AI security becomes default—less friction, fewer surprises.

DIGITAL ERA