PRODIGI
TECHNOLOGY · RISK · TRANSFORMATION

AI GOVERNANCE STARTER FRAMEWORK™

A practical field guide for turning principles into accountable decisions, controls, evidence, and measurable business value.

EXECUTIVE EDITION | VERSION 1.0 | SEPTEMBER 2026

Designed to complement the Prodigi AI Governance Command Center™.

PURPOSE & SCOPE

How to use this framework

Use this guide to establish a minimum viable AI governance program, assess what already exists, and prioritize the controls that matter most. It is intentionally practical: each section names the decision, owner, evidence, and measurement expected.

1. Align leaders

Confirm business objectives, risk appetite, executive sponsor, and who can approve or stop an AI system.

2. Build the inventory

Create one record for each AI system, model, copilot, vendor service, and autonomous agent.

3. Tier the risk

Classify impact, autonomy, data sensitivity, affected groups, and reversibility before deployment.

4. Apply controls

Select safeguards proportionate to risk and connect each control to current evidence.

5. Govern decisions

Record approvals, conditions, exceptions, risk acceptance, and escalation authority.

6. Monitor value

Track drift, incidents, control health, cycle time, adoption, and realized business outcomes.

What this is - and is not

This is an original implementation guide. It does not reproduce any standard and is not legal, audit, certification, or regulatory advice. Organizations should obtain the official source documents and qualified advice for their facts and jurisdictions.

FOUNDATION

Ten operating principles

#	Principle	Practical meaning
01	Business-led	Govern AI in proportion to business impact, not novelty.
02	Named accountability	Every system has an accountable owner and decision authority.
03	Risk-based	Higher impact and autonomy require stronger controls and evidence.
04	Human agency	People retain meaningful authority over consequential outcomes.
05	Purpose-bound	Data, models, tools, and agents operate within approved purpose.
06	Evidence-driven	Claims about safety, performance, and compliance are testable.
07	Traceable	Inputs, decisions, approvals, changes, and incidents are recorded.
08	Secure by design	Threat modeling and access controls start before deployment.
09	Continuous	Material changes trigger reassessment; monitoring continues in use.
10	Value-aware	Governance should enable safe speed and measurable outcomes.

ACCOUNTABILITY

Operating model and decision rights

Role	Primary accountability	Decision right
Executive sponsor	Sets objectives, risk appetite, funding, and escalation expectations.	Accepts material residual risk.
AI governance lead	Runs the program, standards, portfolio view, and governance forum.	Recommends disposition and reports posture.
Business owner	Owns purpose, value, process change, and affected-user outcomes.	Requests use and accepts operational accountability.
Technical owner	Owns architecture, model/vendor controls, testing, security, and change.	Attests technical evidence and readiness.
Risk / Legal / Privacy	Interprets obligations and challenges the control design.	Approves within mandate or escalates.
Internal audit	Provides independent assurance over design and operation.	Does not own first-line controls.

Minimum governance forums

- Portfolio forum: monthly review of inventory, risk, exceptions, incidents, and value.
- Decision forum: approval before high-impact deployment or material change.
- Incident forum: rapid containment, notification, evidence preservation, and lessons learned.

OPERATING PROCESS

The five-stage governance lifecycle

#	Stage	Required work	Output
1	Discover	Register purpose, owner, users, vendor/model, data, integrations, and lifecycle state.	Complete inventory record
2	Assess	Evaluate impact, autonomy, privacy, security, fairness, legal, and operational risks.	Risk tier and assessment
3	Control	Select safeguards, tests, monitoring, human oversight, and evidence requirements.	Control plan and evidence
4	Decide	Approve, approve with conditions, defer, reject, or retire with named accountability.	Decision record
5	Monitor	Track performance, drift, incidents, changes, reviews, controls, cost, and value.	Current posture and actions

Material-change trigger

Re-enter assessment when the system changes purpose, population, model, data source, authority, tool access, geography, vendor, decision impact, or deployment context.

DISCOVER

Minimum AI inventory record

Identity

Unique ID, name, description, lifecycle status, business unit.

Accountability

Executive sponsor, business owner, technical owner, vendor owner.

Purpose & users

Approved purpose, user groups, affected people, prohibited uses.

Technology

Model/provider, version, hosting, integrations, agent tools, dependencies.

Data

Inputs, outputs, sensitive classes, residency, retention, provenance.

Impact

Decision supported, consequence of error, scale, reversibility.

Authority

What the system may read, recommend, draft, execute, or never do.

Governance

Risk tier, approvals, controls, evidence, exceptions, next review.

Monitoring

KPIs, KRIs, thresholds, drift, incidents, change events.

Value

Baseline, target, current result, method, owner, validation status.

ASSESS

A four-tier risk model

Tier	Typical profile	Minimum governance
Tier 1 - Critical	Consequential decisions, safety, rights, essential access, high autonomy, or material financial impact.	Executive approval; independent challenge; strong human authority; intensive testing and monitoring.
Tier 2 - High	Sensitive data, meaningful customer/employee impact, external decisions, or constrained agent execution.	Cross-functional review; documented approval; pre-deployment tests; continuous monitoring.
Tier 3 - Moderate	Internal productivity or decision support with limited impact and reversible outcomes.	Owner approval; standard controls; periodic review; issue escalation.
Tier 4 - Low	Low-impact experimentation or administrative assistance using approved data and tools.	Lightweight registration; acceptable-use controls; owner oversight.

Score the risk - then apply judgment

Consider impact severity, probability, scale, autonomy, data sensitivity, affected groups, transparency, reversibility, external exposure, security, and legal obligations. A numeric score informs the tier; it does not replace accountable judgment.

CONTROL

Control domains

Purpose & accountability

Approved use, owner, users, boundaries, decision rights.

Data governance

Quality, rights, minimization, lineage, retention, access.

Security & resilience

Threat model, secrets, tool permissions, logging, recovery.

Model quality

Fitness, limitations, benchmark design, error analysis, drift.

Fairness & impact

Affected groups, accessibility, bias testing, adverse outcomes.

Transparency

Notice, explainability, provenance, user instructions, disclosures.

Human oversight

Review points, competence, override, appeal, escalation.

Third parties

Due diligence, contracts, change notices, evidence, exit plans.

Agent authority

Action boundaries, approvals, spend limits, kill switch.

Monitoring & response

Thresholds, incidents, reassessment, reporting, remediation.

AGENTIC AI

Agent Authority levels

Level	Permitted authority	Required boundary
1 - Observe	Read approved information and generate summaries.	No external action; no write access.
2 - Recommend	Analyze and propose next actions for human selection.	Recommendations only; source trace required.
3 - Draft	Prepare content, code, tickets, or transactions for approval.	Human approves before release or execution.
4 - Execute	Perform approved, bounded, reversible actions.	Allowlisted tools; limits; monitoring; rapid suspension.
5 - Orchestrate	Coordinate multiple agents/workflows within a defined mandate.	Explicit executive authorization; continuous controls; independent review.

Every agent record should include

- Named owner, approved purpose, authority level, data and tool access.
- Permitted and prohibited actions, approval gates, financial/volume/time limits.
- Logging, review cadence, test evidence, exception handling, and kill-switch owner.

DECIDE

Human oversight and approval gates

Pre-use

Confirm task is within purpose, data is allowed, and user understands limitations.

Pre-action

Require approval before an external, irreversible, high-impact, or out-of-bound action.

Post-action

Review samples, exceptions, outcomes, and tool activity; correct or reverse when possible.

Emergency

Suspend access, contain downstream actions, preserve evidence, notify owners, and investigate.

Approval record

Capture system/agent, requested action, evidence reviewed, risk and conditions, decision, approver, timestamp, rationale, expiration, and required follow-up.

Prohibited without explicit authorization

- Changing its own authority, policies, approval rules, or monitoring.
- Deleting governance evidence, bypassing identity controls, or concealing actions.
- Using unapproved data or tools, transferring funds, or making consequential final decisions.

PROVE

Evidence and audit trail

Evidence area	Minimum artifact
Inventory	Current record, ownership attestation, system diagram.
Assessment	Risk rationale, impact analysis, legal/privacy/security review.
Testing	Test plan, datasets, results, failures, remediation, acceptance.
Controls	Control owner, design, operating evidence, exceptions, expiry.
Decision	Approval, conditions, risk acceptance, rationale, signatures.
Operation	Logs, monitoring results, incidents, complaints, overrides.
Change	Version, data/model/tool changes, reassessment trigger, release.
Value	Baseline, metric definition, source, current result, validation.

Evidence quality test

Evidence should be attributable, current, complete, reproducible, access-controlled, retained for the required period, and connected to the system, control, decision, and version it supports.

MONITOR

Monitoring and incident response

Performance

Accuracy, quality, latency, availability, cost, completion.

Risk

Bias, unsafe output, privacy/security events, user harm, complaints.

Authority

Denied actions, approval bypass attempts, tool errors, spend/volume limits.

Control health

Evidence age, open findings, exceptions, overdue reviews, test failures.

Change

Model/vendor releases, prompt/tool/data updates, integrations, new populations.

Value

Adoption, cycle time, effort avoided, revenue/cost impact, quality outcomes.

Incident response in six actions

Detect the event. Contain access and actions. Assess impact and obligations. Notify the right owners and parties. Remediate root causes and affected outcomes. Learn by updating controls, tests, and authority.

BUSINESS VALUE

Measure governance as a business capability

Dimension	Example measurements
Speed	Governance decision cycle time; percentage within SLA.
Coverage	Percentage of known AI systems inventoried, tiered, and owned.
Control health	Controls passing; evidence current; exceptions overdue.
Agent oversight	Agents with current authority profiles and tested kill switches.
Risk reduction	Critical findings closed; incidents contained; residual risk trend.
Enablement	Initiatives cleared; time-to-safe-deployment; reuse of approved patterns.
Efficiency	Manual hours avoided; assessment reuse; review effort per system.
Outcomes	Quality, revenue, cost, customer, workforce, or mission result validated.

Measurement rule

For every metric record the definition, baseline period, source, owner, target, current result, confidence, validation status, and whether the figure is realized, estimated, or illustrative. Never present a target or pilot estimate as a customer result.

EXTERNAL REFERENCES

Framework alignment overview

Prodigi lifecycle	NIST AI RMF	ISO/IEC 42001	EU AI Act	OWASP agentic guidance
Govern & organize	NIST AI RMF Govern	ISO/IEC 42001 management-system context, leadership, planning and support	Governance, obligations, accountability	Governance of agentic systems
Map & assess	NIST AI RMF Map and Measure	Risk assessment, objectives, operational planning, performance evaluation	Risk classification and applicable requirements	Agentic threat and abuse scenarios
Control & decide	NIST AI RMF Manage	Operational control and improvement	Risk-management duties and human oversight	Identity, tool use, memory, privilege, action boundaries
Monitor & improve	NIST AI RMF Measure and Manage	Monitoring, audit, corrective action, continual improvement	Post-market monitoring and incident duties where applicable	Telemetry, anomaly detection, containment

Alignment note

This high-level crosswalk is a navigation aid, not a statement of certification or legal compliance. Obtain the official publications and perform a requirement-by-requirement assessment for the relevant organization, system, and jurisdiction.

IMPLEMENTATION ROADMAP

A practical 90-day launch plan

Timing	Focus	Key activities	Outputs
Days 1-30	Mobilize	Name sponsor and governance lead; agree scope and risk appetite; inventory priority systems and agents; define tiering and decision forum.	Charter, owner map, initial inventory, risk method
Days 31-60	Baseline	Assess priority systems; define Agent Authority; map controls; collect evidence; identify gaps and quick wins.	Risk register, authority profiles, control/evidence baseline
Days 61-90	Operate	Run approval workflow; remediate priority gaps; establish dashboards, review cadence, incident path, and business-value baseline.	Decision records, action plan, executive scorecard, operating rhythm

Start with a portfolio slice

Choose 5-15 representative systems: one high-impact model, one vendor AI service, one internal copilot, one customer-facing use case, and one autonomous or semi-autonomous agent. This exposes different control needs while keeping the pilot manageable.

NEXT ACTIONS

Executive launch checklist

Ready?	Executive confirmation
☐	An executive sponsor and governance lead are named.
☐	Priority AI systems and agents are inventoried with accountable owners.
☐	Risk tiers and escalation thresholds are approved.
☐	Agent Authority levels, prohibited actions, and kill-switch ownership are documented.
☐	A proportionate control set and evidence standard are defined.
☐	Decision records capture conditions, exceptions, residual risk, and expiry.
☐	Monitoring, incident response, change triggers, and reviews are operational.
☐	Business-value metrics distinguish realized results, estimates, targets, and examples.

Official references

NIST AI Risk Management Framework: https://www.nist.gov/itl/ai-risk-management-framework
NIST AI RMF resources: https://airc.nist.gov/airmf-resources/airmf/
ISO/IEC 42001 overview: https://www.iso.org/standard/42001
European Commission AI regulatory framework: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
OWASP GenAI Security Project: https://genai.owasp.org/

This framework is educational and does not constitute legal, regulatory, audit, certification, cybersecurity, or other professional advice. Adapt it to your organization and obtain qualified advice.

Continue in the Prodigi AI Governance Command Center™

Use the interactive readiness assessment and Agent Authority Toolkit at ai.prodigillc.com/resources.html. For implementation support, engage Prodigi LLC through prodigillc.com.