AI STEWARDSHIP: WHAT DO I DO?

A starting guide for leaders and teams

Start with one real AI endeavor. Establish your Baseline of Stewardship. Act on the gaps.

Why this matters

AI capability is moving into real work faster than many organizations are building the stewardship around it. Governance and controls matter, but stewardship is broader: it includes the human capability, management practices, accountability, monitoring, assurance, resources, learning, maintenance, adaptation, and retirement needed to put AI to work responsibly and productively.

The objective is not maximum control. It is commensurate stewardship — stewardship fit for the capability, use, scale, uncertainty, context, and consequence of the endeavor.

The first move: establish a Baseline of Stewardship

Do not begin by assessing "AI across the whole company." Choose one consequential endeavor — a task, process, service, decision, product, project, or transformation where AI is already being used or seriously considered.

1. START WITH PURPOSE	What are we trying to accomplish? What value matters? What should remain meaningfully human?
2. DEFINE THE ENDEAVOR	What AI-enabled task, process, decision, service, product, or transformation are we actually assessing?
3. UNDERSTAND THE AI CAPABILITY	What can the AI do here? What is uncertain? What data, tools, agents, people, vendors, and other systems does it depend on?
4. DETERMINE REQUIRED STEWARDSHIP	Which stewardship capabilities are applicable, and how strong, integrated, connected, and resourced do they need to be?
5. ESTABLISH THE BASELINE	What stewardship is actually present and working today? What evidence shows that it works in practice?
6. FIND THE MATERIAL GAPS	Where is required stewardship absent, too weak, fragmented, under-scaled, under-resourced, or unable to adapt fast enough?
7. ORGANIZE FOR ACTION	Assign owners. Decide what must be built, strengthened, connected, resourced, restricted, or adapted next.
8. LEARN AND REASSESS	Did the endeavor create the intended value? What changed? What did we learn? Is stewardship still commensurate?

The output is not a score. It is a prioritized stewardship action plan.

© 2026 J L Daniels and Prometheus Endeavor

WHAT SHOULD WE LOOK AT?

Use the questions below as an executive scan — not as a substitute for the full assessment. They are intended to tell you where to look first. Have team members do their own individual scans first. Have them compare with a colleague and revise:

For an individual

- Do I understand what the AI can and cannot reliably do in this situation?
- Am I verifying important outputs in proportion to consequence and uncertainty?
- Am I protecting sensitive information, rights, provenance, and appropriate attribution?
- Am I preserving my own judgment, accountability, agency, and essential skills rather than transferring responsibility to AI?
- Do I know when to seek human expertise, escalate, refuse, correct, or stop?
- Am I learning from outcomes and changing my practice as the tools and consequences change?

With the team members holding a reliable self-image, broaden the scope to the organization or institution:

For an organization or institution

- Direction & accountability — Who owns this endeavor? Who can approve, change, pause, override, or stop it?
- People & competence — Do the people using, supervising, and governing it have the judgment and skills the consequences require?
- Lifecycle controls & assurance — Have we evaluated the system, data, vendors, human oversight, changes, and material risks?
- Operations & learning — Can we detect errors, incidents, drift, misuse, unexpected consequences, and learn from them?
- Security, safety & resilience — What happens when the AI, a vendor, an agent, an integration, or a critical dependency fails or behaves unexpectedly?
- Enterprise integration — Is this still a local experiment, or has AI capability moved to team, cross-functional, or enterprise scale faster than stewardship has?
- Resources & economics — Who funds monitoring, stewardship, maintenance, enhancement, human review, and eventual retirement — not just initial development?
- Human consequences — Who benefits? Who bears consequences? What happens to agency, capability, relationships, trust, opportunity, and human flourishing?

What to do with what you find

Finding	What it means	Default action
Required stewardship is present and working	Protect and sustain it	Monitor, assure, learn, renew
Required stewardship is weak or missing	Material stewardship gap	Build or strengthen; use interim safeguards where needed
AI has scaled beyond stewardship	Scale/integration gap	Strengthen cross-functional ownership, standards,

© 2026 J L Daniels and Prometheus Endeavor

		platforms, controls, learning, and resources
No one can clearly close the gap	Ownership/authority gap	Assign or create the authority, capability, resources, and interfaces needed
Stewardship is stronger than the endeavor requires	Possible overbuild or reusable capacity	Confirm need; redeploy capacity where gaps are more material

Your next 60 minutes

Pick one real AI endeavor. Put the accountable business leader, a knowledgeable user, the relevant technical lead, and the people responsible for risk/assurance or stewardship in the room. Work through the eight steps on page 1. Do not try to solve everything. Identify the three to five stewardship gaps that matter most, name an owner for each, and decide the next action.

Then use the AI Stewardship Assessment Handbook to go deeper.

The Handbook provides the complete Individual, Organization & Institution, and Frontier Development & Public Stewardship normative portfolios and coverage-gap assessment tools.

© 2026 J L Daniels and Prometheus Endeavor