# The AI Stewardship Assessment Handbook

Version 1

October 22, 2026

Address questions or comments to:

Jldaniels6@icloud.com

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Contents and reading routes

Executive Summary ..... 3
Part I Framework and Use ..... 4
Three domains of AI stewardship ..... 5
Start with an endeavor ..... 6
Architecture at a glance ..... 7
Part II Normative Portfolios ..... 8
Individual stewardship 49 elements ..... 9
Organization and Institution stewardship 79 elements ..... 13
Frontier Development and Public Stewardship 76 elements ..... 20
Part III Coverage Gap Assessment Tools ..... 26
Individual assessment ..... 27
Organization and Institution assessment ..... 33
Frontier Development and Public Stewardship assessment ..... 41
Human Flourishing Outcome Review ..... 51
Part IV Cross Domain Review ..... 52
Sample assessment results visualization ..... 53
Part V AI Stewardship What Do I Do ..... 54

## Choose your route

To begin, use Start with an endeavor. For a shorter doorway into the same process, use the two page What Do I Do guide in Part V (page 54).

To assess, select the relevant portfolio in Part II, then use its matching assessment tool in Part III. Read the framework as needed.

To review outcomes and connections, use the Human Flourishing Outcome Review and Part IV Cross Domain Review.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Executive Summary

This field guide helps leaders, AI practitioners, consultants, executives, public officials, and others with stewardship responsibility assess a real AI endeavor, identify material gaps, set priorities, assign responsibility, and decide what to strengthen next.

## AI capability and stewardship

AI capability becomes consequential through use. As individuals, organizations, developers, and public institutions put AI to work, capability can expand through new applications, scale, integration, autonomy, and continuing technical development. AI stewardship is the continuing human responsibility to manage, govern, use, learn from, maintain, assure, adapt, and ultimately retire or replace AI capabilities prudently over their life cycles so benefits can be realized, unacceptable harms reduced, and meaningful human agency preserved. Governance and compliance matter, but stewardship is broader: it is a continuing management responsibility.

The objective is commensurate AI stewardship and sustained human flourishing. Stewardship must be fit for purpose, proportionate to consequence, connected across the relevant actors, and able to learn and adapt as capability and use change.

Capability does not automatically create value. Preserve human learning, independent thinking, and essential skills, and test outcomes over time. When stewardship cannot keep pace with consequential capability, restrict, delay, pause, or decline the activity until appropriate safeguards exist.

## Responsibility and protection

Stewardship operates across individual use, organizations and institutions, and frontier development and public stewardship.

Individuals have stewardship responsibilities that require learned judgment, practice, and appropriate support. These responsibilities must be matched to their role, age, competence, authority, and access to reliable verification.

Individual responsibility does not remove the obligations of developers, providers, deploying institutions, or public authorities. Children and people with limited capacity to recognize or manage AI risks require protection and support that cannot depend on their own verification skills alone.

## Use the handbook

Begin with a defined purpose and endeavor, and pass the preliminary risk and protection gate before experimenting or deploying. The portfolios are repertoires of potentially applicable requirements, rather than universal compliance checklists. Set applicability and required quality before judging current quality; use evidence and consequence to prioritize gaps.

Use the handbook to make stewardship visible and actionable: identify where capability is advancing faster than the capacity to understand, govern, verify, control, learn, resource, and adapt; then close the most consequential gaps. Preserve human agency, strengthen useful human and institutional capability, and keep stewardship learning as AI changes.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Part I Framework and Use

## AI use and capability

Technology creates possibility. Capability is what people and institutions can actually do with technology when they put it to use. As AI use expands, capability can continue to expand through new applications, greater scale, integration into work and institutions, and continued technical development. The stewardship question is therefore not only what AI can do, but what must grow around that capability so it can be used responsibly and productively over time.

## Commensurate stewardship and balanced growth

This handbook uses commensurate stewardship as the unifying assessment idea: stewardship should be fit for purpose and proportionate to capability, use, scale, uncertainty, context, and consequence, and should continuously learn and adapt so that it remains commensurate as those conditions change.

Balanced Growth does not mean that capability and stewardship grow at identical rates or reach a permanent equilibrium. It means that as technological capability and use expand, understanding and learning, effective use and control, and adaptation and renewal grow sufficiently to keep stewardship commensurate. Balanced Growth is iterative: learn, adjust, improve, renew. Relevant stewardship elements must also be sufficiently balanced and connected with one another.

## Stewardship debt

Stewardship Debt is the accumulated gap between the stewardship required to remain commensurate with expanding capability and use and the stewardship actually in place. Debt can arise because needed stewardship was never built, because existing capacity is not strong enough, because new capability creates new demands, or because stewardship capacity erodes.

Paying down the debt means building what is absent, enhancing what is insufficient, maintaining what works, renewing what has weakened, connecting what is fragmented, and adapting as conditions change. Debt includes gaps in learning, judgment, practices, controls, institutions, resources, accountability, and coordination.

## Human flourishing

Value creation is a necessary objective, but it is not the complete outcome frame. Human flourishing asks who benefits, who bears consequences, what human capabilities are strengthened or weakened, and what happens over time. Material dimensions can include material well-being; health and safety; agency and autonomy; capability and competence; meaning and purpose; relationships and social connection; dignity and respect; opportunity and inclusion; distribution of benefits and burdens; freedom from manipulation; privacy and personal sphere; trust and institutional integrity; civic and democratic capacity; resilience and adaptability; and intergenerational sustainability.

These dimensions should not be collapsed into a single universal score. The assessment task is to identify which dimensions are materially affected, establish appropriate evidence and indicators, make boundaries and tradeoffs explicit, assign legitimate decision authority, and monitor effects over time.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

## Three domains of AI stewardship

The practical implication is that responsibility is distributed, but it is not diluted. No single actor can steward the whole AI system. Each actor must steward the part of the system over which it has meaningful capability, authority, influence, or choice, while coordinating with others where consequences cross boundaries.

### Individual Stewardship

Stewardship at the individual level is the responsibility to use AI with informed judgment rather than surrender judgment to it—to understand enough to choose its role, verify what matters, learn from experience, protect the interests of others, and preserve meaningful human agency.

**What is required** Choose when AI is useful and when human attention or expertise should remain primary. Verify important outputs in proportion to consequence and uncertainty. Recognize limitations, incentives, bias, privacy concerns, and the possibility of confident error. Continue learning as tools and uses change rather than relying on fixed assumptions. Remain accountable for one's own choices and avoid treating AI assistance as a transfer of responsibility.

**Defining question** *How do I use AI with judgment, verification, learning, and agency?*

### Organization and Institution Stewardship

Stewardship in organizations and institutions is the responsibility to convert access to AI into reliable, beneficial capability by surrounding its deployment with the governance, skills, controls, monitoring, accountability, and learning needed for the particular mission and context.

**What is required** Decide where AI should and should not be used, and for what purposes. Define human accountability for decisions, outcomes, exceptions, and escalation. Build appropriate policies, skills, security, quality assurance, verification, risk management, and operational controls. Monitor actual behavior and consequences rather than assuming deployment will work as intended. Maintain institutional memory and adapt practices as technology, evidence, and organizational use evolve.

**Defining question** *How do we deploy, govern, monitor, and adapt AI responsibly in the systems for which we are accountable?*

### Frontier Development and Public Stewardship

Stewardship at the frontier is the shared responsibility of AI developers, researchers, governments, standards bodies, and international institutions to shape the conditions under which increasingly capable systems are developed, evaluated, governed, and released

**What is required** Build understanding of emerging capabilities, limitations, and failure modes before and after release. Evaluate and test systems against risks that may extend beyond a single developer or user. Create safeguards, standards, accountability, monitoring, emergency-response capacity, and appropriate release practices. Coordinate where consequences cross firms, sectors, jurisdictions, or national borders. Keep governance adaptive as evidence, capability, and use change.

**Defining question** *How do we develop, evaluate, govern, and release increasingly capable AI responsibly?*

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

## Start with an endeavor

Choose something consequential that a person, organization, developer, or public institution is trying to accomplish. A specific endeavor establishes the capability, context, affected parties, lifecycle, and consequences needed to decide which stewardship elements apply and how strong they must be.

**1. Start with purpose** What are you trying to accomplish? What matters? What should remain meaningfully human?

**2. Choose an endeavor** Pick a real task, problem, process, function, opportunity, research question, service, project, or societal need.

**3. Explore AI capability** What can relevant AI help you do? What model, tool, or combination is appropriate for the endeavor?

**Preliminary risk and protection gate**

Before experimenting or deploying, identify affected parties, potential harms, applicable obligations, uncertainty, and unacceptable boundaries. Decide what safeguards, supervision, testing environment, or restrictions are required before proceeding.

**4. Learn by doing** Prompt, dialogue, iterate, revise, test, verify, and produce. Use real work to develop both AI fluency and human judgment.

**5. Assess stewardship** Given this endeavor, which elements of the relevant normative portfolio are applicable, and how strong must they be?

**6. Close the gaps** Build, enhance, maintain, renew, connect, or adapt the stewardship needed to proceed responsibly.

**7. Evaluate outcomes** Did the endeavor create the intended value? Who benefited? What was strengthened or weakened? What did we learn? Did it contribute to human flourishing?

**8. Iterate** Feed learning back into both capability and stewardship. Adjust the endeavor, the way AI is used, and the stewardship around it as evidence and conditions change.

Learning through real work begins with credible resources and appropriate support. Use practice, verification, reflection, and continuing learning to strengthen both AI fluency and human judgment.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

## The Three Domains of AI Stewardship — Architecture at a Glance

| Stewardship domain | Internal groupings | Elements | Learning responsibility |
|---|---|---|---|
| Individual Stewardship | 1. Understanding & Literacy<br>2. Judgment & Responsible Use<br>3. Safe, Secure & Ethical Practice<br>4. Transparency, Accountability & Recourse<br>5. Monitoring, Feedback & Learning<br>6. Agency, Role Integrity & Human Capability<br>7. Collaboration & Social Stewardship | 49 | Practice responsibly and accountably |
| Organization & Institution Stewardship | 1. Direction, Governance & Accountability<br>2. Knowledge, Capability & Organizational Competence<br>3. Lifecycle Controls, Evaluation & Assurance<br>4. Operational Monitoring, Learning & Adaptation<br>5. People, Rights & Human Consequences<br>6. Security, Safety & Resilience<br>7. External & Institutional Stewardship<br>8. Enterprise Integration, Learning & Value Realization<br>9. Lifecycle Financial & Resource Stewardship | 79 | Enable and govern |
| Frontier Development & Public Stewardship | 1. Frontier Capability Visibility & Shared Understanding<br>2. Evaluation, Testing & Independent Assurance<br>3. Developer Governance & Accountability<br>4. Safety, Security & Control<br>5. Incident Reporting, Learning & Adaptive Stewardship<br>6. Public Governance, Regulation & Institutional Capacity<br>7. Transparency, Public Accountability & Legitimacy<br>8. Ecosystem, Infrastructure & Supply-Chain Stewardship<br>9. International & Cross-Border Stewardship<br>10. Societal Preparedness, Resilience & Public Capability<br>11. Human Flourishing & Societal Direction | 76 | Connect and sustain |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Part II — Normative Portfolios

**Normative Stewardship Portfolio: an organized set of stewardship elements describing the capabilities, practices, responsibilities, and conditions that may be required for commensurate stewardship in a domain.**

The portfolio is a repertoire of possible stewardship requirements, not a checklist that every person, organization, or public system must satisfy in full. The endeavor and its consequences determine what is applicable and how strong stewardship must be.

The normative portfolios describe candidate stewardship coverage. Full concise definitions are provided here. The assessment tools in Part III intentionally do not repeat these definitions; they name the element and move directly to assessment fields.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# A. Individual Stewardship — 49 Elements

Seven groupings organize the 49 elements. Required quality should vary with role and consequence; presence means demonstrated stewardship capacity or behavior, not merely awareness. Apply the portfolio to real human endeavors—skills, tasks, projects, decisions, learning, professional work, and daily activities.

Individual stewardship follows delegated AI activity over its useful life. Selecting or configuring an AI agent or deployment does not transfer accountability to the AI: the responsible person remains accountable for authorized use, supervision, assurance, maintenance and change, interactions with other AI components, continued fitness for purpose, and appropriate discontinuation or disposal.

Individual responsibility does not remove the obligations of developers, providers, deploying institutions, or public authorities. Children and people with limited capacity to recognize or manage AI risks require protection and support that cannot depend on their own verification skills alone.

## 1. Understanding & Literacy

| ID | Stewardship element | Brief definition |
|---|---|---|
| 1 | AI capability literacy | A practical understanding of what relevant AI systems can do, where they are useful, and how their capabilities differ from human capabilities. |
| 2 | Limitations and uncertainty awareness | Recognition that AI outputs can be incomplete, wrong, misleading, context-sensitive, or difficult to explain, and that confidence should be calibrated accordingly. |
| 3 | Context and consequence awareness | Ability to recognize when an AI-assisted action or decision is consequential because it affects safety, rights, opportunity, reputation, resources, or other people. |
| 4 | Data and privacy literacy | Understanding of how prompts, inputs, personal information, proprietary information, and generated content may be collected, retained, exposed, or reused. |
| 5 | Manipulation and synthetic-content literacy | Ability to recognize that AI can generate persuasive, personalized, deceptive, synthetic, or human-like content; to distinguish simulated personhood from human identity or relationship; and to approach uncertain provenance or anthropomorphic framing with appropriate skepticism. |
| 6 | Role-specific stewardship knowledge | Knowledge of the AI-related responsibilities, rules, escalation paths, and professional obligations associated with one's particular role. |
| 7 | Data-quality & provenance judgment | Ability to judge whether data or evidence used as AI input, reference material, training or validation material, or decision support is sufficiently accurate, relevant, representative, current, and traceable for the intended purpose. |

## 2. Judgment & Responsible Use

| ID | Stewardship element | Brief definition |
|---|---|---|
| 8 | Purpose and appropriateness judgment | Ability to decide whether AI is appropriate for the task at all, rather than using it simply because it is available; including recognizing when AI would defeat the human purpose of the task, such as learning, practice, apprenticeship, judgment formation, or productive struggle. |
| 9 | Risk-proportionate use | Ability to increase scrutiny, controls, verification, and human involvement as potential consequence or uncertainty increases. |
| 10 | Human judgment preservation | Maintaining meaningful human judgment rather than treating AI output as a substitute for responsibility, including knowing when to pause, redirect, correct, cancel, seek human review, or stop AI-assisted activity when consequence or uncertainty requires it. |
| 11 | Verification and corroboration | Checking material AI outputs or AI-related claims against reliable evidence, independent sources, domain knowledge, or qualified human review before relying on them; and requiring adequate corroboration before making consequential |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | | judgments about another person's AI use. |
| 12 | Assumption and uncertainty testing | Actively examining the assumptions behind an AI-assisted conclusion and asking what is unknown, ambiguous, or likely to change the answer. |
| 13 | Boundary recognition | Recognizing when a task exceeds one's own competence, authority, or the demonstrated capability of the AI system. |
| 14 | Proportional skepticism | Neither over-trusting nor reflexively rejecting AI or AI-detection claims; calibrating reliance to evidence, context, performance, uncertainty, consequence, and known error patterns. |

## 3. Safe, Secure & Ethical Practice

| ID | Stewardship element | Brief definition |
|---|---|---|
| 15 | Sensitive-information discipline | Avoiding inappropriate entry, disclosure, or exposure of confidential, personal, privileged, security-sensitive, or proprietary information. |
| 16 | Authorized-tool use | Using approved AI systems, accounts, integrations, data sources, agents, and configurations when organizational or professional requirements apply; respecting applicable architecture, interoperability, redundancy, and configuration constraints rather than creating unmanaged or conflicting AI components. |
| 17 | Security hygiene | Following sound security and configuration-management practices for credentials, files, links, plugins, downloads, integrations, AI-generated code or instructions, agent interactions, and material configuration changes. |
| 18 | Misuse and harm avoidance | Refraining from AI uses that are deceptive, coercive, discriminatory, unsafe, unlawful, exploitative, or foreseeably harmful. |
| 19 | Fairness awareness | Considering whether AI-assisted actions, decisions, or enforcement methods could create unjustified differential treatment, exclusion, disadvantage, or false accusation, including through uneven error rates in AI-detection or monitoring tools. |
| 20 | Respect for rights and intellectual property | Considering privacy, consent, attribution, copyright, licensing, ownership, and other relevant rights when using inputs and outputs. |
| 21 | Truthfulness and non-deception | Not presenting AI-generated or AI-assisted material in ways that intentionally misrepresent authorship, evidence, identity, expertise, or certainty. |
| 22 | Consent & permission discipline | Ability to recognize when consent, authorization, participant permission, or another lawful or ethical basis is required before using a person's data, work, likeness, confidential information, or participation in an AI-enabled activity. |
| 23 | Accessibility & equitable-access awareness | Ability to recognize when selection or use of AI creates barriers or unequal access for people with disabilities, limited resources, language differences, or other relevant constraints, and to seek reasonable alternatives or accommodations. |

## 4. Transparency, Accountability & Recourse

| ID | Stewardship element | Brief definition |
|---|---|---|
| 24 | Appropriate AI disclosure | Disclosing material AI involvement when others reasonably need to know it to interpret, trust, review, or act on the work. |
| 25 | Personal accountability | Accepting responsibility for consequential work one submits, approves, communicates, or acts upon even when AI contributed to it. |
| 26 | Decision traceability | Keeping sufficient records of important prompts, sources, assumptions, AI outputs, edits, configurations, changes, evidence, and decisions when later explanation, review, contestation, or consequential attribution of AI use may be needed. |
| 27 | Explainability to others | Being able to explain the basis for an AI-assisted conclusion or action in terms appropriate to the audience rather than relying on 'the AI said so.' |
| 28 | Correction and retraction | Correcting AI-assisted errors promptly and communicating material corrections to people who may have relied on them. |
| 29 | Respect for contestability | Recognizing when another person should be able to question, appeal, request human review, or challenge the evidence supporting an AI-influenced decision or allegation. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| 30 | Source, provenance & citation discipline | Ability to identify and preserve the provenance of important information, distinguish AI-generated assertions from reliable source evidence, and provide attribution or citation when required by professional, academic, legal, or contextual norms. |

## 5. Monitoring, Feedback & Learning

| ID | Stewardship element | Brief definition |
|---|---|---|
| 31 | Outcome monitoring | Attention to whether AI-assisted work actually produces the intended results and to longitudinal effects of repeated AI use on reasoning, confidence, competence, relationships, agency, and human flourishing where material. |
| 32 | Error and anomaly recognition | Noticing unexpected, inconsistent, implausible, biased, unsafe, or otherwise concerning AI behavior. |
| 33 | Incident and near-miss reporting | Reporting significant AI errors, harms, vulnerabilities, misuse, or near misses through appropriate channels. For the endeavor, identify usable reporting channels and who receives, investigates, and acts on reports. Record significant findings, assigned corrective actions, and validation. Correcting an answer or submitting feedback is not evidence that the underlying system has been fixed. |
| 34 | Feedback contribution | Providing useful feedback about system performance, limitations, recurring problems, and effective practices so collective stewardship can improve. |
| 35 | Learning from experience and prudent adaptation | Updating one's practices based on outcomes, errors, new evidence, organizational learning, and changes in AI capability, while using appropriate change control, testing interactions among agents or tools, remaining within authority and applicable safeguards, preserving verification and professional obligations, and accepting accountability for the resulting practice. |
| 36 | Skill renewal | Maintaining relevant AI knowledge and practical competence as systems, risks, policies, and professional expectations evolve. |

## 6. Agency, Role Integrity & Human Capability

| ID | Stewardship element | Brief definition |
|---|---|---|
| 37 | Agency preservation | Using AI in ways that preserve meaningful human choice, responsibility, and the ability to interrupt, redirect, correct, or disengage rather than allowing convenience, dependence, anthropomorphic framing, or automation to displace human agency. |
| 38 | Independent thinking | Retaining the ability and habit of forming one's own view rather than allowing AI to become the default source of framing, interpretation, judgment, or the first response to difficulty; recognizing reflexive dependence and deliberately preserving independent thought. |
| 39 | Core-skill maintenance | Preserving human skills that remain necessary for oversight, fallback, professional competence, independent verification, learning, or development; including deliberate AI-free practice or productive struggle when needed to build or retain those skills. |
| 40 | Appropriate delegation | Distinguishing tasks that can reasonably be delegated to AI from judgments or responsibilities that should remain human; when an individual configures, delegates work to, or relies upon an AI agent, responsibility is not transferred to the agent and includes appropriate selection, authorization, supervision, verification, change awareness, continued fitness for purpose, and discontinuation when no longer appropriate. |
| 41 | Professional and role integrity | Using AI consistently with the duties, standards, boundaries, and legitimate human responsibilities of one's role, including maintaining clear human/AI role boundaries and not treating simulated personhood as a transfer of professional or interpersonal responsibility. |
| 42 | Escalation and refusal | Being willing and able to pause, refuse, or escalate an AI-enabled action when safety, legality, ethics, uncertainty, or authority is in doubt. |
| 43 | Human-flourishing | Ability to recognize material effects of AI use on human flourishing, including well- |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | consequence awareness | being, agency, capability, cognition, relationships, dignity, opportunity, distribution, privacy, trust, resilience, environmental/resource effects, and longer-term consequences. |

## 7. Collaboration & Social Stewardship

| ID | Stewardship element | Brief definition |
|---|---|---|
| 44 | Responsible interaction with others | Considering how one's AI use affects colleagues, customers, students, patients, citizens, family members, or other people who may not have chosen the technology, including effects on trust, mentorship, collaboration, human contact, belonging, and social connection. |
| 45 | Human review and expert consultation | Seeking qualified human input when consequences, ambiguity, specialized knowledge, or uncertainty warrant it. |
| 46 | Shared-practice participation | Following and contributing to team, professional, organizational, or community practices that support responsible AI use. |
| 47 | Constructive challenge | Questioning unsafe assumptions, weak evidence, inappropriate automation, or overconfidence in AI, including when those choices are made by others. |
| 48 | Knowledge sharing | Sharing relevant lessons, effective practices, limitations, and emerging risks so others do not have to relearn them independently. |
| 49 | Collective responsibility awareness | Recognizing that individual AI choices aggregate into organizational and societal consequences and that responsible use is part of a larger stewardship system. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# B. Organization & Institution Stewardship — 79 Elements

Nine groupings organize the 79 elements, including Human Flourishing within the integrated element sequence and lifecycle financial/resource stewardship. Apply the portfolio to the functions, systems, processes, products, services, decisions, and transformations an organization performs today or aspires to improve or reinvent with AI.

## 1. Direction, Governance & Accountability

| ID | Stewardship element | Brief definition |
|---|---|---|
| 1 | Stewardship mandate and objectives | A clear statement of what the organization is responsible for stewarding, the outcomes it seeks to protect or advance, the boundaries of that responsibility, and the shared norms and institutional culture needed to make those commitments meaningful in practice. |
| 2 | Executive and board oversight | Defined senior-level oversight that receives relevant information, challenges decisions, allocates attention, and holds the organization accountable for stewardship performance. |
| 3 | Accountable ownership | Named roles or bodies with explicit responsibility for major stewardship domains, systems, risks, and decisions so responsibility is distributed but not diluted. |
| 4 | Decision rights and escalation | Clear allocation of authority to approve, configure, use, override, interrupt, correct, pause, shut down, or stop AI-enabled activity, including layered model-behavior decision rights among developers/providers, organizational operators, and users where relevant, plus escalation when authority or evidence is insufficient. |
| 5 | Policy and standards framework | An integrated set of policies, standards, procedures, compliance requirements, minimum controls, and shared norms that translates stewardship principles into expected organizational practice while preserving legitimate domain-specific autonomy, professional integrity, academic or scientific freedom, and disciplinary norms where they materially apply. |
| 6 | Risk appetite and boundary conditions | Explicit limits on acceptable risk, prohibited uses, required safeguards, and conditions that trigger heightened review or intervention. |
| 7 | Resource commitment | Sufficient budget, staffing, time, tooling, and authority to perform stewardship responsibilities rather than treating them as unfunded expectations. |

## 2. Knowledge, Capability & Organizational Competence

| ID | Stewardship element | Brief definition |
|---|---|---|
| 8 | Capability and use inventory | A maintained view of relevant AI systems, models, tools, use cases, dependencies, owners, users, and affected processes so the organization knows what it is stewarding. |
| 9 | Technical understanding and expertise | Access to enough technical knowledge to understand system capabilities, limitations, failure modes, dependencies, and uncertainty at the level needed for responsible decisions. |
| 10 | Stewardship competency model | Defined knowledge and skill expectations for leaders, developers, deployers, operators, risk functions, auditors, and other stewardship roles. |
| 11 | Education and role-based training | Ongoing learning that equips people to perform their stewardship responsibilities, recognize emerging issues, and adapt as technology and use change. |
| 12 | Decision-support expertise | Mechanisms to bring legal, risk, security, privacy, ethics, domain, human-factors, and other relevant expertise into consequential decisions. |
| 13 | Institutional memory and knowledge retention | Processes for preserving lessons, rationale, incidents, exceptions, and prior decisions so stewardship learning survives turnover and organizational change. |

## 3. Lifecycle Controls, Evaluation & Assurance

| ID | Stewardship element | Brief definition |
|---|---|---|
| 14 | Use-case intake and | A structured process for identifying proposed uses, their purpose, affected parties, |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | classification | dependencies, and consequence level before deployment or material change. |
| 15 | Risk and impact assessment | A disciplined evaluation of foreseeable harms, benefits, affected stakeholders, misuse, uncertainty, downstream effects, anthropomorphic or relational-design risks, and longitudinal effects on cognition, capability, relationships, agency, and human flourishing, proportionate to the use and its consequences. |
| 16 | Model and system evaluation | Fit-for-purpose testing of performance, robustness, safety, bias or differential effects, security, reliability, and other relevant properties before and during use, including validation of detection or enforcement tools and their false-positive, false-negative, and differential-error consequences before they support consequential action. |
| 17 | Data stewardship | Controls for data quality, provenance, rights, privacy, representativeness, retention, access, and appropriate use across the system lifecycle, including informed consent, authorization, participant rights, withdrawal, confidentiality, and explicit governance of AI interaction logging, retention, audit access, anonymization, and notice where applicable. |
| 18 | Human oversight and intervention design | Deliberate assignment of where human judgment is required, what people must be able to see or understand, and how authorized humans can intervene, redirect, correct, override, pause, cancel, or shut down AI-enabled activity; including protection of learning, professional judgment, employment decision authority, and other human responsibilities that should not be displaced by AI. |
| 19 | Documentation and traceability | Records sufficient to understand what was built or adopted, why decisions were made, what evidence supported them, what changed, who approved key choices, and—where AI interactions are logged or used as evidence—what was retained, how it was obtained, who may access it, and its evidentiary limitations. |
| 20 | Change and release management | Controls for evaluating and approving model, data, prompt, configuration, vendor, workflow, or policy changes that may alter capability, behavior, or risk. |
| 21 | Third-party and supply-chain stewardship | Due diligence and ongoing oversight for vendors, foundation models, data providers, cloud services, contractors, and other external dependencies. |
| 22 | Independent review and assurance | Review by people or functions with enough independence and competence to challenge claims, test controls, validate evidence and detection methods, examine differential error or impact, and provide confidence beyond the operating team's own assessment. |

## 4. Operational Monitoring, Learning & Adaptation

| ID | Stewardship element | Brief definition |
|---|---|---|
| 23 | Operational monitoring | Ongoing observation of system behavior, outcomes, drift, misuse, anomalies, control performance, relevant external changes, and longitudinal human effects—including cognition, capability, relationships, agency, dependence, and human flourishing—plus material financial, environmental, resource, privacy, and logging/auditing consequences after deployment. |
| 24 | Incident identification and reporting | Clear definitions, channels, thresholds, and protections for detecting, documenting, and reporting AI-related incidents, near misses, control failures, unethical or rights-affecting uses, security events, and unexpected consequences. |
| 25 | Incident response and recovery | Prepared procedures, roles, communications, technical actions, and decision authority for containing harm, restoring operations, and protecting affected parties. |
| 26 | Root-cause analysis and corrective action | A disciplined process for learning from incidents and weaknesses, identifying contributing factors, assigning remediation, and verifying that fixes are effective. |
| 27 | Feedback and learning loops | Mechanisms that convert monitoring, incidents, outcomes, user experience, experimentation, and external evidence into responsible organizational learning, changed decisions, improved controls, updated practice, and stewardship adaptation. |
| 28 | Periodic reassessment and renewal | Scheduled and event-triggered review of systems, controls, assumptions, policies, and stewardship capacity to determine whether they remain fit for purpose as |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | | context changes. |

## 5. People, Rights & Human Consequences

| ID | Stewardship element | Brief definition |
|---|---|---|
| 29 | Stakeholder identification and engagement | Systematic identification of people and groups who may benefit, bear risk, hold relevant knowledge, or possess consent, participation, withdrawal, confidentiality, or other rights, with appropriate mechanisms for consultation and participation. |
| 30 | Transparency and communication | Accurate, usable and role-appropriate communication about system purpose, material capabilities and limitations, significant risks, responsibilities, changes, consequential AI involvement, and relevant human/AI role boundaries, including clarity when human-like interaction could be mistaken for personhood or human authority. |
| 31 | User enablement and responsible-use guidance | Instructions, interfaces, constraints, shared norms, and support that help users understand appropriate use, limitations, verification needs, and their own responsibilities; where AI affects learning or professional formation, guidance should clearly communicate permitted and prohibited uses and distinguish the human learning or capability objective from permissible AI assistance. |
| 32 | Contestability, appeal and redress | Ways for affected people to question significant outcomes, seek human review, correct errors, challenge the sufficiency or provenance of AI-related evidence, report harm, and obtain appropriate remedy. |
| 33 | Equity, accessibility and differential-impact review | Attention to whether benefits, burdens, errors, access, exclusion, surveillance, or enforcement fall unevenly across populations, including accessibility for people with differing needs and differential error in AI detection or monitoring tools. |
| 34 | Workforce and role-transition stewardship | Processes for anticipating changes to jobs, skills, authority, workload, professional judgment, apprenticeship, mentorship, novice participation, and developmental opportunities, and for helping people adapt without eroding essential human capability or pathways by which that capability is formed. |

## 6. Security, Safety & Resilience

| ID | Stewardship element | Brief definition |
|---|---|---|
| 35 | Security and misuse protection | Safeguards against unauthorized access, manipulation, model or data theft, prompt or tool abuse, malicious use, and other threats proportionate to system capability. |
| 36 | Resilience and fallback capability | Design and operating arrangements that allow safe degradation, fallback, manual operation, or alternative service when AI systems fail or become unavailable. |
| 37 | Business continuity and crisis integration | Inclusion of AI dependencies and AI-related failure scenarios in business continuity, emergency management, and crisis leadership processes. |
| 38 | Concentration and dependency management | Visibility into critical single points of failure, platform dependencies, vendor concentration, model monocultures, switching barriers, and other systemic dependencies, with deliberate attention to model choice, portability, and avoiding unnecessary lock-in where feasible. |

## 7. External & Institutional Stewardship

| ID | Stewardship element | Brief definition |
|---|---|---|
| 39 | Legal and regulatory alignment | Processes for identifying applicable law, regulation, contractual obligations, and formal duties, and for translating them into accountable operating requirements. |
| 40 | Standards and interoperability participation | Use of relevant technical, professional, assurance, and governance standards, plus participation in standards development when the institution has the capacity or responsibility to contribute. |
| 41 | External incident and risk information sharing | Appropriate mechanisms for sharing significant incidents, vulnerabilities, lessons, and risk signals with regulators, peers, sector bodies, or other institutions where collective learning is needed. |
| 42 | Cross-organizational and | Pre-established relationships, protocols, and decision mechanisms for risks or |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | cross-border coordination | consequences that span organizational, sectoral, jurisdictional, or national boundaries. |
| 43 | Public accountability and reporting | Proportionate disclosure of stewardship commitments, material practices, performance, limitations, incidents, or progress to external stakeholders where legitimacy and trust require it. |
| 44 | Institutional role clarity and gap ownership | Explicit understanding of which stewardship responsibilities belong to the organization, which belong to other institutions, where responsibilities overlap, where no capable owner currently exists, and how shared norms and local/domain-specific authority fit within common institutional stewardship expectations. |
| 45 | Collective capacity building | Contribution to shared stewardship capacity - research, benchmarks, education, incident repositories, assurance methods, public literacy, or other common infrastructure - when individual action is insufficient. |

## 8. Enterprise Integration, Learning & Value Realization

| ID | Stewardship element | Brief definition |
|---|---|---|
| 46 | Enterprise AI strategy and value thesis | A coherent enterprise view of where AI can create material organizational value, how that value supports strategy, and what capabilities and organizational changes are required to realize it. |
| 47 | Enterprise use-case portfolio management | A disciplined mechanism for discovering, comparing, prioritizing, funding, scaling, stopping, and learning from AI use cases across the organization. |
| 48 | Process and workflow redesign capability | The organizational capability to redesign end-to-end work around AI rather than simply adding AI tools to existing tasks, while preserving or deliberately replacing the human learning, apprenticeship, mentorship, relationship, judgment, and professional-development functions embedded in existing work. |
| 49 | Cross-functional transformation orchestration | Governance and operating mechanisms for coordinating AI-enabled changes that span functions, processes, data, systems, roles, and decision rights. |
| 50 | Enterprise platforms, standards and reusable capabilities | Common AI platforms, tools, architectures, data practices, security patterns, controls, and reusable components that allow successful approaches to scale safely, including protected institutional AI environments for sensitive, restricted, confidential, or regulated information and, where feasible, model-agnostic access that preserves model choice and reduces lock-in. |
| 51 | Value measurement and benefits realization | Methods for establishing baselines, measuring productivity and process outcomes, tracking financial and non-financial value, and determining whether AI investments produce sustained enterprise benefit, including material effects on human capability, relationships, access, environmental/resource footprint, and other stakeholder outcomes. |
| 52 | Learning-to-scale mechanism | A deliberate institutional mechanism for enabling, governing, and capturing learning from individual and team experimentation; evaluating evidence; standardizing what works; propagating effective practices; and updating controls, skills, decision rights, and operating practices as AI capability and use change. |
| 53 | Operating-model adaptation | The ability to change roles, structures, workflows, decision rights, management systems, and interfaces as AI changes how work is performed, while deliberately preserving or rebuilding apprenticeship, mentorship, collaboration, professional formation, and other human-development pathways that may otherwise be displaced. |
| 54 | Incentive and performance alignment | Incentives and performance measures that encourage enterprise outcomes, responsible experimentation, shared learning, and cross-functional optimization rather than local gains alone. |
| 55 | Executive transformation governance | Active executive direction of AI-enabled enterprise transformation, including strategic priorities, cross-functional dependencies, resource allocation, metrics, and repeated decisions about the next increment of change. |
| 56 | Scale-proportionate | A mechanism ensuring governance, controls, assurance, monitoring, skills, and |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | stewardship architecture | institutional capacity become stronger and more integrated as AI moves from individual experimentation to teams, processes, and enterprise-wide operations. |
| 57 | Enterprise integration and interoperability | Technical and operating integration that connects AI-enabled processes, data, systems, models, and organizational units without creating unmanaged dependencies or fragmented local solutions, and that supports portability or model choice where these reduce concentration or lock-in risk. |
| 58 | Transformation pace management | Deliberate calibration of the rate of AI adoption and organizational change to the enterprise's capacity to learn, integrate, govern, resource, and absorb change without allowing transformation to outrun stewardship. |
| 59 | Transformation sequencing and dependency management | Deliberate ordering of AI-enabled changes so necessary capabilities, infrastructure, controls, skills, process understanding, and organizational prerequisites are established before dependent transformations proceed. |
| 60 | Endeavor chunking and learning increments | Structuring AI-enabled transformation into manageable, consequential increments that produce usable outcomes and evidence, permit learning and course correction, and avoid committing the enterprise too far ahead of what it has learned. |
| 61 | Organizational absorptive capacity | The enterprise's ability to absorb simultaneous AI-enabled changes without overwhelming leadership attention, workforce capacity, technical integration, controls, learning mechanisms, or operational stability. |
| 62 | Human Flourishing Impact Assessment & Measurement | A disciplined process for identifying, measuring, and monitoring how AI-enabled changes affect material dimensions of human flourishing, including well-being, health and safety, agency, cognition, competence, meaning, relationships, dignity, opportunity, distribution, freedom from manipulation, privacy, trust, civic capacity, resilience, apprenticeship and mentorship opportunities, environmental/resource effects, and longitudinal consequences where material. |
| 63 | Flourishing Tradeoff & Boundary Governance | A governance process for identifying conflicts between value creation and human flourishing; establishing unacceptable boundaries; making material tradeoffs explicitly; considering who benefits and who bears consequences; protecting meaningful human capabilities, agency, relationships, apprenticeship and developmental pathways; and assigning authority and accountability for those decisions. |

## 9. Lifecycle Financial & Resource Stewardship

| ID | Stewardship element | Brief definition |
|---|---|---|
| 64 | Full-lifecycle AI cost model | Estimate the complete resource requirement from development/acquisition/deployment through Monitoring & Stewardship (MS), Maintenance (M), Enhancement (E), and retirement/replacement. |
| 65 | Monitoring & Stewardship (MS) funding | Durably fund the people, tools, assurance, monitoring, review, incident response, governance, learning, and stewardship adaptation required while the capability remains in use. |
| 66 | Maintenance funding | Sustain funding for technical, data, integration, security, vendor, documentation, infrastructure, and operational maintenance. |
| 67 | Enhancement funding | Fund adaptation and improvement as capability, use, risk, regulation, workflows, and organizational needs change. |
| 68 | Development-to-operating-cost transition | Define the financial and organizational drop-down from temporary development/project funding into recurring production, MS, maintenance, and enhancement budgets. |
| 69 | Lifecycle ownership and budget transfer | Explicitly transfer operational responsibility and recurring budget when an experiment, pilot, or development endeavor becomes an enduring capability. |
| 70 | Stewardship liability recognition | Recognize at approval that deployment creates future MS+M+E, assurance, learning, security, governance, and retirement obligations. |
| 71 | Lifecycle value-to-cost assessment | Compare expected sustained value with full lifecycle cost, including MS+M+E and retirement, rather than original development cost alone, and include material environmental/resource costs and human-development opportunity costs where |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | | they affect the stewardship decision. |
| 72 | Stewardship reserve and funding sufficiency | Determine whether future recurring stewardship obligations are adequately funded or provisioned rather than accumulating as deferred stewardship. |
| 73 | Portfolio stewardship burden management | Measure and manage the cumulative recurring MS+M+E burden created by the installed AI portfolio. |
| 74 | Funding priority and tradeoff governance | Allocate scarce lifecycle resources according to consequence, value, risk, dependency, and organizational priorities rather than automatically favoring new development. |
| 75 | Retirement and decommissioning economics | Plan and fund migration, archival obligations, dependency removal, data disposition, replacement, and decommissioning. |
| 76 | Vendor and external-service lifecycle economics | Account for recurring model/API charges, data costs, vendor changes, assurance, switching costs, external dependencies, portability, model-choice requirements, and the lifecycle economics of avoiding or exiting provider lock-in. |
| 77 | Human stewardship labor accounting | Recognize the human time required for prompting, verification, judgment, review, escalation, correction, monitoring, documentation, learning, and oversight. |
| 78 | Lifecycle cost and value reforecasting | Periodically update expected MS+M+E costs, retirement obligations, risks, realized value, vendor/model dependencies, and material environmental/resource effects using operating evidence. |
| 79 | Financial exit and stop criteria | Define economic and stewardship conditions for pausing, redesigning, replacing, or retiring a capability. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# C. Frontier Development & Public Stewardship — 76 Elements

Eleven groupings organize the 76 elements. Frontier stewardship is a distributed system; assessment must consider both element quality and whether information, authority, expertise, accountability, and action are sufficiently connected. Apply the portfolio to the capabilities, public purposes, institutions, infrastructure, services, and societal or community outcomes that frontier developers and public actors seek to create, preserve, improve, or govern.

## 1. Frontier Capability Visibility & Shared Understanding

| ID | Stewardship element | Brief definition |
|---|---|---|
| 1 | Frontier capability monitoring | Ongoing observation of significant advances in frontier AI capabilities, methods, scale, autonomy, and deployment potential so stewardship can respond to material changes. |
| 2 | Common capability measurement | Shared benchmarks, evaluation methods, and terminology that allow developers, governments, researchers, and other institutions to compare consequential capabilities meaningfully. |
| 3 | Frontier threshold definition | Criteria for identifying when models, systems, training runs, or capabilities cross thresholds that warrant heightened stewardship obligations. |
| 4 | Emerging-risk identification | Structured processes for identifying new hazards, misuse pathways, control problems, and societal consequences before they become routine or widespread. |
| 5 | Uncertainty characterization | Explicit treatment of what is known, unknown, contested, or difficult to measure about frontier capabilities, behavior, and consequences. |
| 6 | Public-interest research capacity | Independent research capability sufficient to understand frontier systems, evaluate claims, investigate risks, and inform public stewardship rather than relying exclusively on developers. |

## 2. Evaluation, Testing & Independent Assurance

| ID | Stewardship element | Brief definition |
|---|---|---|
| 7 | Pre-deployment frontier evaluations | Rigorous evaluation of consequential capabilities, failure modes, misuse potential, autonomy, controllability, and safeguards before frontier systems are broadly released or deployed. |
| 8 | Standardized high-consequence evaluations | Common evaluation protocols for capabilities with potentially severe consequences, such as cyber, biological, chemical, deception, autonomous action, or other dangerous applications. |
| 9 | Independent testing access | Mechanisms allowing qualified independent evaluators appropriate access to frontier systems, evidence, or secure testing environments. |
| 10 | Red-team and adversarial testing | Systematic efforts to discover vulnerabilities, dangerous behaviors, safeguard bypasses, and unexpected interactions under realistic adversarial conditions, including creative abuse brainstorming and scenario generation that draws on diverse technical and nontechnical perspectives. |
| 11 | Evaluation integrity and reproducibility | Controls that make important evaluation results sufficiently reliable, documented, repeatable, and resistant to selective reporting or benchmark gaming. |
| 12 | Third-party assurance | Independent assessment of whether frontier developers' claimed safeguards, controls, evaluations, and governance practices are actually operating as represented. |
| 13 | Post-deployment evaluation | Continued testing after release to detect capability changes, emergent uses, distribution shifts, new failure modes, and weaknesses not evident before deployment. |

## 3. Developer Governance & Accountability

| ID | Stewardship element | Brief definition |
|---|---|---|
| 14 | Frontier safety and | Internal governance structures at frontier developers, including an appropriate |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | stewardship governance | internal governing body or equivalent mechanism, with clear authority over safety, security, deployment, and stewardship decisions. |
| 15 | Senior leadership accountability | Named executive accountability for frontier stewardship commitments, risk decisions, and corrective action. |
| 16 | Governing-body oversight | Board or equivalent oversight of material frontier risks, major capability thresholds, deployment decisions, and stewardship capacity. |
| 17 | Risk thresholds and decision rules | Explicit criteria linking measured capabilities or risks to required safeguards, deployment restrictions, escalation, or non-deployment. |
| 18 | Internal challenge and independence | Functions with sufficient independence, expertise, information, and authority to challenge capability-development and deployment decisions. |
| 19 | Whistleblower and protected escalation | Safe mechanisms for employees and contractors to raise material safety, security, legal, or public-interest concerns without retaliation. |
| 20 | Stewardship commitment transparency | Public clarity about major frontier stewardship commitments, decision frameworks, material changes, and performance against those commitments. |

## 4. Safety, Security & Control

| ID | Stewardship element | Brief definition |
|---|---|---|
| 21 | Frontier model security | Protection of model weights, training infrastructure, sensitive research, credentials, and other assets whose theft or compromise could create substantial harm. |
| 22 | Secure development infrastructure | Security controls proportionate to the sensitivity and potential consequence of frontier training, experimentation, deployment, and supporting systems. |
| 23 | Dangerous-capability safeguards | Technical and procedural controls designed to prevent, restrict, or detect access to capabilities that could enable severe misuse. |
| 24 | Access and release controls | Mechanisms to calibrate who can access frontier capabilities, under what conditions, with what monitoring, and at what level of system openness. |
| 25 | Alignment and controllability research | Sustained research into whether increasingly capable systems can be directed, constrained, monitored, interrupted, and kept aligned with intended objectives. |
| 26 | Human control and intervention | Practical ability for authorized humans or institutions to monitor and intervene, restrict, pause, or terminate consequential system behavior or deployment, with explicit attention to whether human detection and response can operate fast enough for the capability and timescale involved. |
| 27 | Emergency response and containment | Prepared procedures, authority, coordination, and technical means to respond rapidly enough to severe frontier-system incidents or newly discovered hazards, including automated or pre-authorized safeguards where human response time alone may be insufficient. |
| 28 | Resilience and recovery | Plans and capabilities to restore safe operation, contain propagation, learn from failure, and protect critical functions after a frontier-related disruption, including scenarios in which system actions may unfold faster than ordinary human oversight. |

## 5. Incident Reporting, Learning & Adaptive Stewardship

For the endeavor, identify usable reporting channels and who receives, investigates, and acts on reports. Record significant findings, assigned corrective actions, and validation. Correcting an answer or submitting feedback is not evidence that the underlying system has been fixed.

| ID | Stewardship element | Brief definition |
|---|---|---|
| 29 | Mandatory material-incident reporting | Requirements or strong mechanisms for timely reporting of significant frontier AI incidents, near misses, security breaches, dangerous capabilities, and control failures, including accessible routes for individual and organizational users to submit material reports. |
| 30 | Shared incident taxonomy | Common definitions and classification schemes, aligned where appropriate with established or newly developed technical and governance standards, that make incidents comparable across developers, jurisdictions, and sectors. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| 31 | Confidential reporting channels | Usable mechanisms for sharing sensitive incident or vulnerability information with appropriate authorities or trusted institutions without forcing unsafe public disclosure, with clear recipients, escalation routes, and responsibility for action. |
| 32 | Incident investigation capacity | Independent or public-interest capability to investigate serious AI incidents, determine causes, and recommend corrective action. |
| 33 | Cross-system learning and adaptation | Mechanisms that allow evidence, evaluation results, incidents, near misses, corrective actions, and effective stewardship practices to improve decisions across developers, independent evaluators, public institutions, sectors, and—where appropriate—jurisdictions, so learning can propagate through the distributed stewardship system. |
| 34 | Corrective-action tracking | Processes to ensure significant findings lead to assigned actions, deadlines, validation, and closure rather than remaining advisory. |
| 35 | Periodic stewardship reassessment | Regular and event-triggered reassessment of safeguards, rules, institutions, assumptions, and stewardship capacity as frontier capability, evidence, deployment, use, and consequences change. |

## 6. Public Governance, Regulation & Institutional Capacity

| ID | Stewardship element | Brief definition |
|---|---|---|
| 36 | Competent public oversight institutions | Public bodies with sufficient technical expertise, authority, resources, independence, and speed to oversee frontier development credibly. |
| 37 | Regulatory visibility | Mechanisms giving appropriate public authorities sufficient information about frontier training, capabilities, evaluations, deployments, and material incidents. |
| 38 | Proportionate frontier requirements | Rules that increase stewardship obligations with capability, scale, risk, or consequence rather than treating all AI development identically. |
| 39 | Licensing, registration, or authorization mechanisms | Where warranted by consequence, mechanisms for requiring notification, registration, licensing, or approval of specified frontier activities. |
| 40 | Inspection and information-access authority | Authority for competent oversight bodies to obtain evidence, conduct or commission evaluations, and verify compliance. |
| 41 | Enforcement and corrective authority | Credible powers and processes to require remediation, impose restrictions, halt unsafe activity, or apply proportionate sanctions. |
| 42 | Regulatory learning and adaptation | Processes that allow public governance to learn from evidence, evaluations, incidents, societal outcomes, and implementation experience and to update rules, thresholds, expertise, institutions, and practices as conditions change. |
| 43 | Institutional independence and conflict safeguards | Protections against regulatory capture, inappropriate political or commercial influence, and conflicts that undermine public-interest oversight. |

## 7. Transparency, Public Accountability & Legitimacy

| ID | Stewardship element | Brief definition |
|---|---|---|
| 44 | Frontier-system transparency | Appropriate disclosure of consequential capabilities, limitations, evaluations, safeguards, and deployment conditions without creating avoidable security risks. |
| 45 | Risk and safety reporting | Periodic public reporting on material risks, incidents, evaluation findings, safety investments, and stewardship performance. |
| 46 | Traceability and documentation | Records sufficient to reconstruct consequential development, evaluation, release, and governance decisions when review is required, with tamper-evident or tamper-resistant logging and, where warranted, append-only or write-once/read-many storage that prevents systems or operators from silently altering the record. |
| 47 | Public-interest scrutiny | Meaningful opportunities for qualified researchers, civil society, professional bodies, and other public-interest actors to examine frontier stewardship. |
| 48 | Stakeholder and affected-party engagement | Mechanisms to hear from communities and groups likely to bear significant benefits, risks, or externalities from frontier AI. |
| 49 | Decision legitimacy and due process | Clear, reviewable processes for consequential public or developer decisions, including appropriate reasons, appeal, challenge, or independent review. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| 50 | Protection of legitimate transparency | Legal and institutional safeguards for good-faith research, reporting, journalism, and disclosure that contribute to public stewardship. |

## 8. Ecosystem, Infrastructure & Supply-Chain Stewardship

| ID | Stewardship element | Brief definition |
|---|---|---|
| 51 | Compute and critical-infrastructure visibility | Appropriate visibility into concentrations and uses of advanced compute and other infrastructure relevant to frontier development and systemic risk. |
| 52 | Cloud and infrastructure-provider stewardship | Defined responsibilities for major compute, cloud, and infrastructure providers that enable frontier training or deployment. |
| 53 | Supply-chain security | Controls addressing vulnerabilities in hardware, software, data, model components, vendors, and services supporting frontier systems. |
| 54 | Third-party and downstream risk management | Stewardship of how frontier models are integrated, fine-tuned, redistributed, or used by downstream actors. |
| 55 | Research and capability diffusion management | Processes for balancing scientific openness and beneficial diffusion against the risk of disseminating information or capabilities that materially enable severe harm. |
| 56 | Concentration and systemic-dependency monitoring | Assessment of risks created by reliance on a small number of frontier models, developers, infrastructure providers, or technical components. |

## 9. International & Cross-Border Stewardship

| ID | Stewardship element | Brief definition |
|---|---|---|
| 57 | International capability and risk dialogue | Sustained channels among governments, developers, scientists, and institutions for shared understanding of frontier capabilities and risks. |
| 58 | Interoperable frontier standards | Technical and governance standards that can operate across jurisdictions and reduce dangerous gaps or incompatible requirements. |
| 59 | Cross-border incident notification | Arrangements for rapid notification and coordination when frontier AI incidents or vulnerabilities may affect multiple countries. |
| 60 | International evaluation cooperation | Mechanisms for sharing methods, evidence, expertise, and where appropriate evaluation results across trusted national and international institutions. |
| 61 | Coordination on extreme-risk safeguards | Agreements or compatible approaches for capabilities whose consequences could extend beyond the jurisdiction in which they are developed. |
| 62 | Regulatory-gap and jurisdiction-arbitrage management | Cooperation intended to reduce incentives or opportunities to relocate high-risk activity solely to evade meaningful stewardship. |
| 63 | Global participation and capacity building | Efforts to ensure countries with less technical or institutional capacity can participate in and benefit from frontier stewardship rather than remaining outside the system. |

## 10. Societal Preparedness, Resilience & Public Capability

| ID | Stewardship element | Brief definition |
|---|---|---|
| 64 | Public capability literacy | Accessible public understanding of consequential frontier AI capabilities, limitations, uncertainty, and realistic implications. |
| 65 | Professional and sector preparedness | Preparation of critical professions and sectors to recognize, manage, and adapt to frontier-enabled changes and risks, including use or development of appropriate professional standards, competency expectations, training, certification, or licensing mechanisms where justified by consequence and professional context. |
| 66 | Critical-infrastructure preparedness | Assessment and strengthening of critical systems that may be affected by advanced AI-enabled cyber, information, operational, or other threats. |
| 67 | Information-environment resilience | Capacity to detect, understand, and respond to large-scale synthetic media, manipulation, impersonation, fear amplification, hype, coordinated persuasion, or other intentionally or problematically amplified AI-enabled information threats while |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship element | Brief definition |
|---|---|---|
| | | preserving legitimate debate and expression. |
| 68 | Economic and workforce transition capacity | Institutions and policies capable of monitoring and responding to significant labor-market, organizational, or economic changes driven by frontier AI. |
| 69 | Public emergency preparedness | Plans for low-probability but high-consequence frontier AI events requiring coordinated public response. |
| 70 | Societal impact monitoring | Ongoing monitoring of material societal consequences as frontier AI diffuses, including distributional, institutional, democratic, economic, security, workforce, cognitive, capability, relationship, dependency, and human-flourishing effects over time. |

## 11. Human Flourishing & Societal Direction

| ID | Stewardship element | Brief definition |
|---|---|---|
| 71 | Human-flourishing impact assessment & measurement | A disciplined process for identifying, measuring, and monitoring how frontier AI development, release, diffusion, and public-policy choices affect material dimensions of human flourishing, including well-being, health and safety, agency, cognition, competence, meaning, relationships, dignity, opportunity, distribution, freedom from manipulation, privacy, trust, civic capacity, resilience, environmental/resource effects, and intergenerational sustainability. |
| 72 | Flourishing tradeoff & boundary governance | Publicly legitimate processes for identifying conflicts between capability or economic value and human flourishing; establishing unacceptable boundaries; making material tradeoffs explicitly; considering who benefits and who bears consequences; protecting meaningful human capabilities and agency; and assigning authority and accountability for those decisions. |
| 73 | Human-capability and agency preservation | Assessment of whether frontier AI strengthens or erodes human judgment, learning, competence, autonomy, meaningful choice, professional capability, dignity and esteem, material security, purpose, and the ability of people and institutions to function without excessive dependence on AI. |
| 74 | Distributional benefit and burden stewardship | Mechanisms for identifying how frontier-AI benefits, costs, displacement, risks, access, power, and opportunities are distributed across populations, sectors, regions, countries, and generations, with attention to groups that may have limited voice or capacity to protect their interests. |
| 75 | Social connection, meaning and institutional integrity | Monitoring and stewardship of effects on human relationships, social participation, meaningful roles, trust, institutional legitimacy, and the social practices through which people and communities develop capability, belonging, purpose, and mutual obligation. |
| 76 | Long-term and intergenerational stewardship | Processes for evaluating cumulative, path-dependent, difficult-to-reverse, environmental, resource, institutional, and intergenerational consequences of frontier AI, including whether near-term gains create longer-term dependency, fragility, capability loss, or burdens transferred to future populations. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Part III — Coverage-Gap Assessment Tools

**Assessment Instrument: a learning and development tool for determining which portfolio elements apply to a particular endeavor, how strong they need to be, what evidence demonstrates current stewardship, and what must be learned, built, enhanced, connected, or adapted next.**

The instruments are not intended to assess AI adoption for its own sake. They begin with a real endeavor and help identify the stewardship needed to use or develop AI responsibly and successfully. Assessment should support learning, action, and improvement—not become an end in itself.

**Microsoft-derived assessment evidence.** Where relevant, assessment should test whether people and institutions can exercise practical human control; distinguish human-like AI behavior from human identity, relationship, or authority; detect longitudinal effects of repeated AI use on cognition, capability, relationships, dependence, and flourishing; and understand who holds model-behavior decision rights at each layer of development, operation, and use.

Required Quality means the level at which an applicable stewardship element must be designed, implemented, resourced, practiced, evidenced, effective, and adaptable to be fit for purpose. It is the required quality of the stewardship capability represented by the element—not a second meaning of AI capability.

The assessment instruments use the portfolio element names without repeating portfolio definitions. For full definitions, refer to Part II. Worksheet header rows are set to repeat automatically on subsequent pages.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# A. Individual Stewardship Coverage-Gap Assessment — 49 Elements

## Purpose

This instrument converts the revised 49-element Individual stewardship portfolio into a structured assessment of whether a person has the stewardship capacity appropriate to a defined role and AI-use context. The central distinction remains that presence means demonstrated capacity or behavior—not merely awareness, policy acknowledgment, or completion of training.

## Assessment Unit and Context

| Assessment field | Entry |
|---|---|
| Individual / role | |
| Business, professional, civic, or personal context | |
| AI systems or capability in scope | |
| Typical AI-assisted tasks / decisions | |
| Affected parties | |
| Level of autonomy / decision authority | |
| Potential consequence if AI is misused or wrong | |
| Assessment date / assessor | |

## Six-Step Assessment Logic

**1. Applicability by role and context —** Determine whether the element is required, potentially relevant, or not applicable.

**2. Presence / demonstrated capability —** Determine whether the individual can actually demonstrate the element in realistic situations.

**3. Quality / fit for purpose —** Assess how reliably, independently, consistently, and appropriately the person applies the element.

**4. Consequence weighting —** Estimate consequence if the individual lacks or fails to apply the element when needed.

**5. Gap prioritization —** Compare required quality with demonstrated quality and combine the gap with consequence.

**6. Appropriate response —** Select remediation that addresses the cause and context: education, practice, coaching, changed decision rights, stronger human review, technical guardrails, restricted use, or organizational-system change.

## Applicability Anchors

| Rating | Meaning | Use |
|---|---|---|
| A — Applicable | Materially required for this role/context. | Set required quality and assess demonstrated quality. |
| P — Potentially Applicable | Depends on unresolved facts, future use, or conditional responsibilities. | Resolve before final scoring; consider interim safeguards. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Rating | Meaning | Use |
|---|---|---|
| N — Not Applicable | Not materially required in the defined role/context. | Document rationale and reassess after material change. |

## Demonstrated Quality / Fit-for-Purpose Anchors

| Score | Anchor | Demonstrated capability | Illustrative evidence |
|---|---|---|---|
| 0 | Not demonstrated | Cannot demonstrate the element when required, or behavior contradicts it. | Scenario failure, observed unsafe practice, inability to explain or perform required action. |
| 1 | Aware / prompted | Recognizes or performs only with substantial prompting, supervision, or correction. | Knowledge check plus coached exercise; limited observation. |
| 2 | Developing / supported | Can apply in routine situations with guidance, but judgment is inconsistent in novel or higher-consequence situations. | Scenario exercises, supervised practice, sampled work. |
| 3 | Proficient / independent | Reliably recognizes and applies independently across the normal range required by the role, including escalation. | Observed work, realistic simulations, review of decisions and outcomes. |
| 4 | Adaptive / exemplary | Applies reliably in novel/difficult situations, calibrates to consequence, learns from outcomes, and strengthens practice around them. | Complex simulations, sustained operating evidence, peer contribution, improvement evidence. |

## Evidence Hierarchy

| Evidence strength | Examples | Interpretation |
|---|---|---|
| Strong | Observed behavior in real work; realistic scenario/simulation; documented decisions/outcomes; demonstrated escalation or correction. | Best evidence that the element is actually present and usable. |
| Moderate | Work samples; structured interview with concrete examples; supervisor/peer evidence; repeated practical exercises. | Useful corroboration when direct observation is limited. |
| Weak | Training completion; policy acknowledgment; self-rating; simple knowledge quiz. | Evidence of exposure or awareness, not by itself evidence of stewardship capacity. |

## Consequence Weighting

| Weight | Anchor | If the element is absent or fails when needed |
|---|---|---|

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Weight | Anchor | If the element is absent or fails when needed |
|---|---|---|
| 1 | Low | Limited, localized, readily reversible effects; little impact on others. |
| 2 | Moderate | Meaningful effect on work quality, privacy, resources, reputation, colleagues, customers, or other affected people. |
| 3 | High | Serious safety, rights, financial, professional, legal, employment, or other consequential effects. |
| 4 | Critical | Severe, widespread, difficult-to-reverse, life/safety, systemic, or mission-critical consequences. |

## Gap and Priority Rule

For each Applicable element, set Required Quality (1–4) before rating Demonstrated Quality (0–4). Gap Size = max(Required Quality − Demonstrated Quality, 0). Priority Score = Gap Size × Consequence Weight. The score is a decision aid, not a substitute for judgment.

| Score | Priority | Default response |
|---|---|---|
| 0 | No current gap | Maintain through normal use, feedback, and periodic reassessment. |
| 1–2 | Low | Targeted learning or practice; normal supervision. |
| 3–4 | Moderate | Formal development action, coaching, added review, or workflow support; reassess. |
| 5–8 | High | Prompt remediation; stronger supervision/human review or guardrails until capability improves. |
| 9–12 | Very High | Management/professional oversight; change decision rights or restrict relevant AI use. |
| 13–16 | Critical | Immediate protective action; pause/prohibit consequential activity unless adequate compensating controls are in place. |

### 1. Understanding & Literacy

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 1 | AI capability literacy | | | | | | |
| 2 | Limitations and uncertainty awareness | | | | | | |
| 3 | Context and consequence awareness | | | | | | |
| 4 | Data and privacy literacy | | | | | | |
| 5 | Manipulation and synthetic-content | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| | literacy | | | | | | |
| 6 | Role-specific stewardship knowledge | | | | | | |
| 7 | Data-quality & provenance judgment | | | | | | |

## 2. Judgment & Responsible Use

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 8 | Purpose and appropriateness judgment | | | | | | |
| 9 | Risk-proportionate use | | | | | | |
| 10 | Human judgment preservation | | | | | | |
| 11 | Verification and corroboration | | | | | | |
| 12 | Assumption and uncertainty testing | | | | | | |
| 13 | Boundary recognition | | | | | | |
| 14 | Proportional skepticism | | | | | | |

## 3. Safe, Secure & Ethical Practice

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 15 | Sensitive-information discipline | | | | | | |
| 16 | Authorized-tool use | | | | | | |
| 17 | Security hygiene | | | | | | |
| 18 | Misuse and harm avoidance | | | | | | |
| 19 | Fairness awareness | | | | | | |
| 20 | Respect for rights and intellectual property | | | | | | |
| 21 | Truthfulness and non-deception | | | | | | |
| 22 | Consent & permission discipline | | | | | | |
| 23 | Accessibility & equitable-access awareness | | | | | | |

## 4. Transparency, Accountability & Recourse

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 24 | Appropriate AI disclosure | | | | | | |
| 25 | Personal accountability | | | | | | |
| 26 | Decision traceability | | | | | | |
| 27 | Explainability to others | | | | | | |
| 28 | Correction and retraction | | | | | | |
| 29 | Respect for contestability | | | | | | |
| 30 | Source, provenance & citation discipline | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

## 5. Monitoring, Feedback & Learning

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 31 | Outcome monitoring | | | | | | |
| 32 | Error and anomaly recognition | | | | | | |
| 33 | Incident and near-miss reporting | | | | | | |
| 34 | Feedback contribution | | | | | | |
| 35 | Learning from experience and prudent adaptation | | | | | | |
| 36 | Skill renewal | | | | | | |

## 6. Agency, Role Integrity & Human Capability

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 37 | Agency preservation | | | | | | |
| 38 | Independent thinking | | | | | | |
| 39 | Core-skill maintenance | | | | | | |
| 40 | Appropriate delegation | | | | | | |
| 41 | Professional and role integrity | | | | | | |
| 42 | Escalation and refusal | | | | | | |
| 43 | Human-flourishing consequence awareness | | | | | | |

## 7. Collaboration & Social Stewardship

| ID | Stewardship Element | App. | Req. Q | Demo Q | Conseq. | Evidence / Demonstration | Response / Owner / Due |
|---|---|---|---|---|---|---|---|
| 44 | Responsible interaction with others | | | | | | |
| 45 | Human review and expert consultation | | | | | | |
| 46 | Shared-practice participation | | | | | | |
| 47 | Constructive challenge | | | | | | |
| 48 | Knowledge sharing | | | | | | |
| 49 | Collective responsibility awareness | | | | | | |

## Individual gap remediation record

| Field | Assessment / action |
|---|---|
| **Element / role context** | |
| **Why applicable** | |
| **Required vs. demonstrated quality** | |
| **Consequence and rationale** | |
| **Evidence** | |
| **Root cause / development need** | |
| **Response / remediation** | |
| **Interim safeguards** | |
| **Owner / coach** | |
| **Target date** | |
| **Validation method** | |
| **Reassessment / residual gap** | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# B. Organization & Institution Coverage-Gap Assessment — 79 Elements

## Purpose

This instrument assesses the complete 79-element, nine-domain Organization & Institution stewardship portfolio for applicability, presence, fit for purpose, scale/integration, transformation dynamics, lifecycle resource sufficiency, human-flourishing consequences, priority, ownership, and remediation.

## Required Lifecycle Resource Model

### Development / Acquisition / Deployment → MS + M + E → Retirement / Replacement

| Component | Resource obligation |
|---|---|
| Development / Acquisition / Deployment | Build/buy, integration, data preparation, testing, change, training, implementation, and launch. |
| MS — Monitoring & Stewardship | Monitoring, evaluation, human review, assurance, governance, incidents, compliance, learning, consequence observation, and stewardship adaptation. |
| M — Maintenance | Technical, model, prompt, data, integration, security, infrastructure, vendor, documentation, and operational upkeep. |
| E — Enhancement | Improvements and adaptations as capability, use, risks, regulation, workflows, and needs change. |
| Retirement / Replacement | Migration, records, dependency removal, contract exit, data disposition, replacement, and decommissioning. |

Establish an initial lifecycle estimate, identify uncertainty, assign operational and financial ownership, fund the obligation, and reforecast from actual operating evidence. Do not assume a universal AI MS+M+E ratio.

## Integrated Assessment Logic

**1. Applicability and presence —** Determine whether the element is required and whether a recognizable capability or mechanism exists.

**2. Quality / fit for purpose —** Rate design, implementation, resourcing, evidence, effectiveness, and adaptability.

**3. Scale / integration fit —** Determine whether stewardship operates at the organizational level at which capability and expected value reside.

**4. Transformation-dynamics fit —** Assess pace, sequencing, chunking, and absorptive capacity.

**5. Lifecycle resource fit —** Determine whether MS+M+E and retirement obligations are estimated, assigned, funded, sustainable, and reforecast.

**6. Consequence and priority —** Weight weakness or absence by consequence and prioritize the largest material gap.

**7. Ownership —** Assign accountable operational and financial owners, including budget acceptance at production transition.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

**8. Remediation and learning —** Close gaps, validate improvement, update lifecycle cost/value assumptions, and feed evidence into subsequent decisions.

## Integrated Rating Anchors

| Score | Quality / fit | Scale / integration | Lifecycle resource fit |
|---|---|---|---|
| 0 | Absent | Isolated / absent | No lifecycle model, owner, or recurring funding. |
| 1 | Initial / ad hoc | Individual / local | Some recurring costs recognized informally; funding is reactive or temporary. |
| 2 | Defined | Team / point process | MS+M+E categories, owners, and estimates defined, but funding, evidence, or transition is incomplete. |
| 3 | Operational | Cross-functional / scaled | Recurring obligations are budgeted, owned, monitored, and reforecast using operating evidence. |
| 4 | Adaptive / high quality | Enterprise / adaptive | Lifecycle economics are integrated into portfolio decisions; costs, value, reserves, retirement, and stewardship capacity are continuously rebalanced. |

### Connection to the Emerging Plateau Concept

The existing Scale / Integration Anchors provide the assessment bridge to the emerging Plateau concept. They already distinguish Isolated / absent, Individual / local, Team / point process, Cross-functional / scaled, and Enterprise / adaptive operation. The emerging Plateau concept uses a closely related progression to describe how AI capability and organizational leverage expand from individual use through team/department use, cross-functional leverage, and enterprise-wide integration. For this Working Draft 1.9, the original Scale / Integration Anchors are restored unchanged. The Plateau terminology is not substituted for, or used to rewrite, these anchors; their relationship should be tested before any later nomenclature change.

## Priority Rule

Set Required Quality (RQ), Required Scale/Integration (RS), and where relevant Required Lifecycle Resource Fit (RF). Quality Gap=max(RQ−CQ,0); Scale Gap=max(RS−CS,0); Financial/Resource Gap=max(RF−CF,0). For transformation-dynamics elements also assess the transformation gap. Priority Score = largest applicable gap × Consequence Weight (1–4).

## AI Endeavor Lifecycle Resource Worksheet

| Lifecycle field | Estimate / owner / evidence |
|---|---|
| Endeavor / capability | |
| Business owner | |
| Technical owner | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Lifecycle field | Estimate / owner / evidence |
|---|---|
| Stewardship owner | |
| Financial owner | |
| Original development / acquisition / deployment cost | |
| Expected operating life | |
| Annual MS cost | |
| Annual Maintenance cost | |
| Annual Enhancement cost | |
| Expected retirement / replacement cost | |
| Human stewardship labor requirement | |
| External/vendor recurring cost | |
| Infrastructure/data recurring cost | |
| Total expected lifecycle stewardship burden | |
| Expected annual / lifecycle value | |
| Key cost and value uncertainties | |
| Development-to-production budget transfer date | |
| Operational owner accepting MS+M+E obligation | |
| Funding source / cost center | |
| Reserve / contingency | |
| Reforecast frequency | |
| Economic stop / retirement criteria | |

## 1. Direction, Governance & Accountability

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 1 | Stewardship mandate and objectives | | | | | | |
| 2 | Executive and board oversight | | | | | | |
| 3 | Accountable ownership | | | | | | |
| 4 | Decision rights and escalation | | | | | | |
| 5 | Policy and standards framework | | | | | | |
| 6 | Risk appetite and boundary conditions | | | | | | |
| 7 | Resource commitment | | | | | | |

## 2. Knowledge, Capability & Organizational Competence

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 8 | Capability and use inventory | | | | | | |
| 9 | Technical understanding and expertise | | | | | | |
| 10 | Stewardship competency model | | | | | | |
| 11 | Education and role-based training | | | | | | |
| 12 | Decision-support expertise | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 13 | Institutional memory and knowledge retention | | | | | | |

## 3. Lifecycle Controls, Evaluation & Assurance

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 14 | Use-case intake and classification | | | | | | |
| 15 | Risk and impact assessment | | | | | | |
| 16 | Model and system evaluation | | | | | | |
| 17 | Data stewardship | | | | | | |
| 18 | Human oversight and intervention design | | | | | | |
| 19 | Documentation and traceability | | | | | | |
| 20 | Change and release management | | | | | | |
| 21 | Third-party and supply-chain stewardship | | | | | | |
| 22 | Independent review and assurance | | | | | | |

## 4. Operational Monitoring, Learning & Adaptation

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 23 | Operational monitoring | | | | | | |
| 24 | Incident identification and reporting | | | | | | |
| 25 | Incident response and recovery | | | | | | |
| 26 | Root-cause analysis and corrective action | | | | | | |
| 27 | Feedback and learning loops | | | | | | |
| 28 | Periodic reassessment and renewal | | | | | | |

## 5. People, Rights & Human Consequences

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 29 | Stakeholder identification and engagement | | | | | | |
| 30 | Transparency and communication | | | | | | |
| 31 | User enablement and responsible-use guidance | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 32 | Contestability, appeal and redress | | | | | | |
| 33 | Equity, accessibility and differential-impact review | | | | | | |
| 34 | Workforce and role-transition stewardship | | | | | | |

## 6. Security, Safety & Resilience

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 35 | Security and misuse protection | | | | | | |
| 36 | Resilience and fallback capability | | | | | | |
| 37 | Business continuity and crisis integration | | | | | | |
| 38 | Concentration and dependency management | | | | | | |

## 7. External & Institutional Stewardship

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 39 | Legal and regulatory alignment | | | | | | |
| 40 | Standards and interoperability participation | | | | | | |
| 41 | External incident and risk information sharing | | | | | | |
| 42 | Cross-organizational and cross-border coordination | | | | | | |
| 43 | Public accountability and reporting | | | | | | |
| 44 | Institutional role clarity and gap ownership | | | | | | |
| 45 | Collective capacity building | | | | | | |

## 8. Enterprise Integration, Learning & Value Realization

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 46 | Enterprise AI strategy and value thesis | | | | | | |
| 47 | Enterprise use-case portfolio management | | | | | | |
| 48 | Process and workflow redesign capability | | | | | | |
| 49 | Cross-functional transformation | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| | orchestration | | | | | | |
| 50 | Enterprise platforms, standards and reusable capabilities | | | | | | |
| 51 | Value measurement and benefits realization | | | | | | |
| 52 | Learning-to-scale mechanism | | | | | | |
| 53 | Operating-model adaptation | | | | | | |
| 54 | Incentive and performance alignment | | | | | | |
| 55 | Executive transformation governance | | | | | | |
| 56 | Scale-proportionate stewardship architecture | | | | | | |
| 57 | Enterprise integration and interoperability | | | | | | |
| 58 | Transformation pace management | | | | | | |
| 59 | Transformation sequencing and dependency management | | | | | | |
| 60 | Endeavor chunking and learning increments | | | | | | |
| 61 | Organizational absorptive capacity | | | | | | |
| 62 | Human Flourishing Impact Assessment & Measurement | | | | | | |
| 63 | Flourishing Tradeoff & Boundary Governance | | | | | | |

## 9. Lifecycle Financial & Resource Stewardship

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| 64 | Full-lifecycle AI cost model | | | | | | |
| 65 | Monitoring & Stewardship (MS) funding | | | | | | |
| 66 | Maintenance funding | | | | | | |
| 67 | Enhancement funding | | | | | | |
| 68 | Development-to-operating-cost transition | | | | | | |
| 69 | Lifecycle ownership | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Scale | Req./Current Resource Fit | Consequence / Priority | Owners / Gap / Remediation |
|---|---|---|---|---|---|---|---|
| | and budget transfer | | | | | | |
| 70 | Stewardship liability recognition | | | | | | |
| 71 | Lifecycle value-to-cost assessment | | | | | | |
| 72 | Stewardship reserve and funding sufficiency | | | | | | |
| 73 | Portfolio stewardship burden management | | | | | | |
| 74 | Funding priority and tradeoff governance | | | | | | |
| 75 | Retirement and decommissioning economics | | | | | | |
| 76 | Vendor and external-service lifecycle economics | | | | | | |
| 77 | Human stewardship labor accounting | | | | | | |
| 78 | Lifecycle cost and value reforecasting | | | | | | |
| 79 | Financial exit and stop criteria | | | | | | |

## Organization & Institution management-level review

- Where are the largest concentrations of material stewardship gaps?
- Where has AI deployment scaled faster than stewardship capacity?
- Which successful experiments have no mechanism for scaling or institutional learning?
- Where do local standards, platforms, data practices, controls, or incentives prevent enterprise reuse?
- Which cross-functional dependencies lack an accountable orchestrator?
- Are lifecycle monitoring, maintenance, enhancement, and retirement obligations owned and funded?
- Which human-flourishing dimensions are materially affected, and who has authority over tradeoffs and boundaries?
- Where is stewardship no longer commensurate with capability, use, scale, uncertainty, or consequence?

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# C. Frontier Development & Public Stewardship Coverage-Gap Assessment — 76 Elements

## Purpose

This instrument converts the Frontier Development & Public Stewardship candidate portfolio into a structured method for identifying and evaluating stewardship coverage gaps across a distributed public-private, national, and international system.

It applies the common Hidden AI Challenge logic: determine applicability; identify presence; assess quality or fit for purpose; weight the consequence of weakness or absence; prioritize the gap; assign responsibility; and define remediation.

This domain adds two system-level tests: who can actually close the gap when authority and capability are distributed, and whether the relevant elements are connected well enough that individually present components combine into effective stewardship coverage.

## Assessment Scope

| Assessment field | Definition / entry |
|---|---|
| Assessment boundary | |
| Frontier capability / activity in scope | |
| Developers / providers in scope | |
| Public authorities / jurisdictions | |
| Independent evaluators / research bodies | |
| Infrastructure / ecosystem actors | |
| International bodies / arrangements | |
| Affected populations / sectors | |
| Assessment date / lead assessor | |

## Seven-Step Element Assessment

**1. Applicability -** Determine whether the element is materially required for the capability, activity, actors, jurisdictions, or consequences in scope.

**2. Presence -** Determine whether a recognizable mechanism, institution, authority, practice, capability, or arrangement actually exists.

**3. Quality / fit for purpose -** Assess mandate, competence, capacity, independence, authority, evidence, reach, speed, and adaptability.

**4. Consequence weighting -** Estimate consequence if the element is absent, weak, fragmented, bypassed, or fails when needed.

**5. Gap prioritization -** Compare required quality with current quality and combine the gap with consequence.

**6. Responsibility / closability -** Identify the actor or combination of actors with the authority, capability, legitimacy, and leverage to close the gap.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

**7. Remediation -** Define the intervention required, including strengthening, connecting, or creating institutional capacity.

## Applicability Anchors

| Rating | Meaning | Treatment |
|---|---|---|
| A - Applicable | Materially required within the assessment boundary. | Set required quality and complete full assessment. |
| P - Potentially Applicable | Relevance depends on unresolved capability, use, jurisdictional, or consequence questions. | Resolve before final scoring; consider interim coverage. |
| N - Not Applicable | Not materially required within the defined boundary. | Document rationale and reassess after material change. |

## Quality / Fit-for-Purpose Anchors

| Score | Anchor | System meaning | Illustrative characteristics |
|---|---|---|---|
| 0 | Absent | No meaningful element or coverage exists. | No responsible mechanism; no reliable capacity or evidence. |
| 1 | Initial / Fragmented | Some activity exists but is partial, weak, voluntary, limited, or bypassable. | Patchy coverage; unclear authority; limited independence or reach. |
| 2 | Defined / Developing | Formal mandate, roles, processes, or institutions exist but implementation or capacity is incomplete. | Material operating, coordination, enforcement, or coverage gaps. |
| 3 | Operational / Fit for purpose | Functioning with adequate authority, expertise, resources, independence, reach, evidence, coordination, and responsiveness. | Routine operation; credible evidence; effective interfaces. |
| 4 | Adaptive / High quality | Robust under changing frontier conditions, independently challenged, connected, evidence-driven, and able to renew quickly. | Anticipatory capacity; tested coordination; continuous learning. |

## System Connectivity Anchors

Rate connectivity separately from element quality. A strong element can still leave a systemic gap if information, authority, escalation, or action does not connect across the stewardship system.

| Score | Anchor | Meaning |
|---|---|---|
| 0 | Disconnected | Relevant actors or elements do not reliably exchange information, coordinate decisions, or act together. |
| 1 | Ad hoc | Connections depend on personal relationships, voluntary contact, or crisis improvisation. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Score | Anchor | Meaning |
|---|---|---|
| 2 | Defined | Interfaces, agreements, or channels exist, but use, coverage, speed, or authority is inconsistent. |
| 3 | Operational | Information, escalation, coordination, and joint action work reliably under normal and stressed conditions. |
| 4 | Adaptive / resilient | Connections are tested, cross-border or cross-institutional where needed, learn from events, and adjust as capability and risk change. |

## Consequence Weighting

| Weight | Anchor | Potential consequence |
|---|---|---|
| 1 | Low | Limited, localized, readily reversible consequences. |
| 2 | Moderate | Meaningful sectoral, institutional, economic, rights, security, or public-trust consequences. |
| 3 | High | Serious cross-organizational or cross-jurisdictional harm, major disruption, substantial rights/safety impacts, or systemic exposure. |
| 4 | Critical | Severe, widespread, transnational, difficult-to-reverse, catastrophic, or civilization-scale potential consequences. |

## Gap Priority Rule

For Applicable elements: Element Gap = max(Required Quality - Current Quality, 0). Connectivity Gap = max(Required Connectivity - Current Connectivity, 0). Base Priority = max(Element Gap, Connectivity Gap) x Consequence Weight. The higher gap is used because either a weak element or a broken interface can defeat coverage.

Elevate priority when the gap is time-sensitive, hard to detect, spans jurisdictions, involves concentrated private capability, lacks any actor with clear authority, or interacts with other gaps.

| Score | Priority | Default response |
|---|---|---|
| 0 | No current gap | Maintain, test, monitor, and reassess. |
| 1-2 | Low | Normal improvement cycle; clarify minor interfaces. |
| 3-4 | Moderate | Assign formal remediation, owner(s), milestones, and validation. |
| 5-8 | High | Prompt cross-actor action; strengthen authority, capacity, information flow, or safeguards. |
| 9-12 | Very High | Senior public/private leadership attention; consider restrictions or compensating measures until reduced. |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Score | Priority | Default response |
|---|---|---|
| 13-16 | Critical | Immediate coordinated action; consider pause, prohibition, emergency authority, or extraordinary safeguards where appropriate. |

## Responsibility / Closability Assessment

A gap is not truly assigned if the named actor lacks the authority, capability, legitimacy, jurisdiction, information, resources, or cooperation required to close it.

| Question | Assessment | Illustrative actors |
|---|---|---|
| Who has authority? | | Developer; regulator; legislature; standards body; infrastructure provider; treaty body. |
| Who has technical capability? | | Developer; evaluator; research institution; security agency; specialist lab. |
| Who has information? | | Developer; cloud provider; incident reporter; regulator; affected sector. |
| Who has legitimacy / mandate? | | Public authority; professional body; international institution; independent assurance body. |
| Who can act at required scale? | | National government; coordinated regulators; major providers; international coalition. |
| Who must cooperate? | | Multiple developers; public-private actors; jurisdictions; infrastructure providers. |
| Is there an orphan gap? | | No actor currently possesses enough authority/capability to close it alone; new institutional capacity may be required. |

## Remediation Menu

| Remediation type | Illustrative actions |
|---|---|
| Strengthen an existing element | Increase expertise, funding, authority, independence, evaluation quality, enforcement, security, monitoring, or adaptive capacity. |
| Clarify responsibility or authority | Assign statutory, contractual, governance, or operational responsibility; define decision rights and escalation. |
| Create a missing institution or function | Establish an evaluator, incident body, regulator, coordinating office, research capability, assurance function, or other missing capacity. |
| Connect existing elements | Create reporting channels, interfaces, joint protocols, escalation routes, shared taxonomies, exercises, or interoperable systems. |
| Establish information-sharing arrangements | Create secure reporting, confidential disclosure, trusted exchange, or mandatory notification mechanisms. |
| Improve independence and challenge | Separate assurance from commercial incentives; protect |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| Remediation type | Illustrative actions |
|---|---|
| | internal challenge; fund independent public-interest expertise. |
| Align standards and thresholds | Develop common definitions, evaluation methods, capability thresholds, assurance criteria, or reciprocal recognition. |
| Build cross-border coordination | Create bilateral, plurilateral, or international arrangements for evaluation, incidents, safeguards, enforcement, or extreme-risk response. |
| Apply interim safeguards | Restrict access, deployment, scaling, release, or specific capabilities until adequate stewardship coverage exists. |
| Build societal preparedness | Strengthen sector readiness, critical infrastructure, public literacy, workforce transition, emergency response, or resilience. |

## 1. Frontier Capability Visibility & Shared Understanding

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 1 | Frontier capability monitoring | | | | | | |
| 2 | Common capability measurement | | | | | | |
| 3 | Frontier threshold definition | | | | | | |
| 4 | Emerging-risk identification | | | | | | |
| 5 | Uncertainty characterization | | | | | | |
| 6 | Public-interest research capacity | | | | | | |

## 2. Evaluation, Testing & Independent Assurance

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 7 | Pre-deployment frontier evaluations | | | | | | |
| 8 | Standardized high-consequence evaluations | | | | | | |
| 9 | Independent testing access | | | | | | |
| 10 | Red-team and adversarial testing | | | | | | |
| 11 | Evaluation integrity and reproducibility | | | | | | |
| 12 | Third-party assurance | | | | | | |
| 13 | Post-deployment evaluation | | | | | | |

## 3. Developer Governance & Accountability

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 14 | Frontier safety and stewardship governance | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 15 | Senior leadership accountability | | | | | | |
| 16 | Governing-body oversight | | | | | | |
| 17 | Risk thresholds and decision rules | | | | | | |
| 18 | Internal challenge and independence | | | | | | |
| 19 | Whistleblower and protected escalation | | | | | | |
| 20 | Stewardship commitment transparency | | | | | | |

## 4. Safety, Security & Control

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 21 | Frontier model security | | | | | | |
| 22 | Secure development infrastructure | | | | | | |
| 23 | Dangerous-capability safeguards | | | | | | |
| 24 | Access and release controls | | | | | | |
| 25 | Alignment and controllability research | | | | | | |
| 26 | Human control and intervention | | | | | | |
| 27 | Emergency response and containment | | | | | | |
| 28 | Resilience and recovery | | | | | | |

## 5. Incident Reporting, Learning & Adaptive Stewardship

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 29 | Mandatory material-incident reporting | | | | | | |
| 30 | Shared incident taxonomy | | | | | | |
| 31 | Confidential reporting channels | | | | | | |
| 32 | Incident investigation capacity | | | | | | |
| 33 | Cross-system learning and adaptation | | | | | | |
| 34 | Corrective-action tracking | | | | | | |
| 35 | Periodic stewardship reassessment | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

## 6. Public Governance, Regulation & Institutional Capacity

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 36 | Competent public oversight institutions | | | | | | |
| 37 | Regulatory visibility | | | | | | |
| 38 | Proportionate frontier requirements | | | | | | |
| 39 | Licensing, registration, or authorization mechanisms | | | | | | |
| 40 | Inspection and information-access authority | | | | | | |
| 41 | Enforcement and corrective authority | | | | | | |
| 42 | Regulatory learning and adaptation | | | | | | |
| 43 | Institutional independence and conflict safeguards | | | | | | |

## 7. Transparency, Public Accountability & Legitimacy

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 44 | Frontier-system transparency | | | | | | |
| 45 | Risk and safety reporting | | | | | | |
| 46 | Traceability and documentation | | | | | | |
| 47 | Public-interest scrutiny | | | | | | |
| 48 | Stakeholder and affected-party engagement | | | | | | |
| 49 | Decision legitimacy and due process | | | | | | |
| 50 | Protection of legitimate transparency | | | | | | |

## 8. Ecosystem, Infrastructure & Supply-Chain Stewardship

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 51 | Compute and critical-infrastructure visibility | | | | | | |
| 52 | Cloud and infrastructure-provider stewardship | | | | | | |
| 53 | Supply-chain security | | | | | | |
| 54 | Third-party and downstream risk management | | | | | | |
| 55 | Research and | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| | capability diffusion management | | | | | | |
| 56 | Concentration and systemic-dependency monitoring | | | | | | |

## 9. International & Cross-Border Stewardship

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 57 | International capability and risk dialogue | | | | | | |
| 58 | Interoperable frontier standards | | | | | | |
| 59 | Cross-border incident notification | | | | | | |
| 60 | International evaluation cooperation | | | | | | |
| 61 | Coordination on extreme-risk safeguards | | | | | | |
| 62 | Regulatory-gap and jurisdiction-arbitrage management | | | | | | |
| 63 | Global participation and capacity building | | | | | | |

## 10. Societal Preparedness, Resilience & Public Capability

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 64 | Public capability literacy | | | | | | |
| 65 | Professional and sector preparedness | | | | | | |
| 66 | Critical-infrastructure preparedness | | | | | | |
| 67 | Information-environment resilience | | | | | | |
| 68 | Economic and workforce transition capacity | | | | | | |
| 69 | Public emergency preparedness | | | | | | |
| 70 | Societal impact monitoring | | | | | | |

## 11. Human Flourishing & Societal Direction

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 71 | Human-flourishing impact assessment & measurement | | | | | | |

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

| ID | Stewardship Element | App. | Req./Current Quality | Req./Current Connectivity | Consequence / Priority | Responsible Actor / Gap Closer | Remediation / Status |
|---|---|---|---|---|---|---|---|
| 72 | Flourishing tradeoff & boundary governance | | | | | | |
| 73 | Human-capability and agency preservation | | | | | | |
| 74 | Distributional benefit and burden stewardship | | | | | | |
| 75 | Social connection, meaning and institutional integrity | | | | | | |
| 76 | Long-term and intergenerational stewardship | | | | | | |

## Frontier system-level coverage review

- Where are the largest concentrations of High, Very High, or Critical gaps?
- Which gaps have no actor with sufficient authority and capability to close them?
- Where do strong elements fail because interfaces or information flows are weak?
- Where can actors bypass stewardship by changing provider, jurisdiction, release model, or infrastructure?
- Which gaps require cross-border coordination?
- Which elements depend excessively on voluntary action or a single institution?
- Where is independent challenge insufficient?
- Can the system detect and respond to rapid capability change?
- Are incident learning and corrective action connected across developers, public institutions, sectors, and jurisdictions?
- What new institution, authority, standard, or coordination mechanism would close multiple gaps at once?

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Human Flourishing Outcome Review

For consequential AI endeavors, assess human-flourishing consequences alongside value, risk, security, and stewardship capacity. This is not a single composite score. Identify material dimensions, affected populations, indicators, boundaries, tradeoffs, decision authority, trends, and corrective action.

| Flourishing dimension | Material? | Finding / trend / evidence |
|---|---|---|
| **Material well-being** | | |
| **Health & safety** | | |
| **Agency & autonomy** | | |
| **Capability & competence** | | |
| **Meaning & purpose** | | |
| **Relationships & social connection** | | |
| **Dignity & respect** | | |
| **Opportunity & inclusion** | | |
| **Distribution of benefits & burdens** | | |
| **Freedom from manipulation** | | |
| **Privacy & personal sphere** | | |
| **Trust & institutional integrity** | | |
| **Civic & democratic capacity** | | |
| **Resilience & adaptability** | | |
| **Intergenerational sustainability** | | |

## Four decision questions

- Who benefits?
- Who bears the consequences?
- What human capabilities are strengthened or weakened?
- What happens over time?

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Part IV — Cross-Domain Review

A complete stewardship review may require more than one domain. Individual capacity can be strong while institutional support is weak; an organization can have strong internal controls while depending on missing public or frontier capacity; public rules can exist while organizations and individuals lack the competence or systems to implement them. Review interfaces among domains, not merely scores within them.

| Cross-domain question | Finding / evidence |
|---|---|
| **Where does capability live, and does stewardship follow it?** | |
| **Which gaps originate in another stewardship domain?** | |
| **Where is responsibility distributed, overlapping, or unclear?** | |
| **Can evidence and learning move across the relevant boundaries fast enough?** | |
| **Where is stewardship no longer commensurate with capability, use, scale, uncertainty, context, or consequence?** | |
| **What stewardship debt is accumulating?** | |
| **Which human-flourishing dimensions are materially affected?** | |
| **Which of the four responsibilities—Mobilize, Build & Enhance, Embed, Direct—requires the most attention now?** | |

## Interpretive Principle — Commensurate Stewardship

The purpose of assessment is not to maximize stewardship activity. It is to determine whether stewardship is commensurate with the capability and use in scope: fit for purpose, proportionate to consequence, connected at the required scale, adequately resourced, and able to learn and adapt as conditions change. A gap exists when needed stewardship is absent, too weak, poorly connected, under-resourced, or unable to adapt fast enough.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Sample Assessment-Results Visualization

After completing an element-level assessment, practitioners can roll results up into a simple two-dimensional management view. Plot each applicable stewardship element by Current Quality on the horizontal axis and Required Quality on the vertical axis. Consequence or priority can be shown alongside the element. This makes material gaps visible without replacing the underlying assessment.

| Position | Interpretation | Typical response |
|---|---|---|
| High required / Low current | Priority stewardship gap | Build or strengthen promptly; consider interim safeguards. |
| High required / High current | Protect and sustain | Maintain, monitor, assure, learn, and renew. |
| Low required / Low current | Monitor or defer | Do not overbuild; revisit as capability, use, or consequence changes. |
| Low required / High current | Potential overinvestment or transferable capacity | Confirm the requirement; consider redeploying capacity where gaps are more material. |

## Illustrative 2×2 — Required Quality × Current Quality

| | LOW CURRENT QUALITY | HIGH CURRENT QUALITY |
|---|---|---|
| HIGH REQUIRED QUALITY | PRIORITY GAP<br>Strengthen | PROTECT & SUSTAIN<br>Maintain / assure |
| LOW REQUIRED QUALITY | MONITOR / DEFER<br>Avoid overbuilding | REVIEW / REDEPLOY<br>Potential excess capacity |

Interpretation. Use the 2×2 as a management view, not as a substitute for element-level evidence. A high-consequence element may warrant priority even when its numerical gap is modest. Organization & Institution and Frontier/Public assessments should also consider scale, lifecycle resource fit, connectivity, dependencies, and whether the responsible actor has the authority and capability to close the gap.

#### Selected references and further reading

A concise selected-reference appendix may be added in a later release for readers who want to trace the stewardship, institutional, human-flourishing, and standards foundations without interrupting the handbook's practitioner flow.

#### Sources and Development Basis

Working Draft 1.9 consolidates the Hidden AI Challenge stewardship framework, the three-domain stewardship definitions, the current normative portfolios, and the corresponding coverage-gap assessment methods into a single reader-facing repository. The conceptual framing draws from The Hidden AI Challenge integrated presentation and performance materials; the stewardship definitions draw from Stewardship Across the Three AI Domains; and the detailed repository uses the current Individual, Organization & Institution, and Frontier Development & Public Stewardship portfolios and assessment instruments. University of Florida and MIT materials informed prior refinements that are now integrated into the relevant elements rather than presented as separate additions. Microsoft AI's 2026 Humanist AI Code of Conduct informed the Working Draft 1.8 refinements to human control, human/AI role boundaries, longitudinal human-impact evaluation, and layered model-behavior decision rights.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# Part V AI Stewardship What Do I Do

*A two-page starting guide for leaders and teams | Guide v0.2*

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

## Why this matters

AI is moving into real work faster than many organizations are building the stewardship it requires: human capability, management, accountability, controls, monitoring, assurance, resources, learning, maintenance, adaptation, and retirement.

The objective is commensurate stewardship: fit for the endeavor's capability, use, scale, uncertainty, context, and consequence, so AI contributes to human flourishing.

Capability does not automatically create value. Preserve human learning, independent thinking, and essential skills, and test outcomes over time. When stewardship cannot keep pace with consequential capability, restrict, delay, pause, or decline the activity until appropriate safeguards exist.

## The first move: establish a Baseline of Stewardship

Choose one consequential AI endeavor: a task, process, service, decision, product, project, or transformation already in use or under consideration.

Before experimenting or deploying, identify affected parties, potential harms, applicable obligations, uncertainty, and unacceptable boundaries. Decide what safeguards, supervision, testing environment, or restrictions are required before proceeding.

| | |
|---|---|
| **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 and next actions. Build, strengthen, connect, resource, restrict, or adapt what is needed. |
| **8. LEARN AND REASSESS** | Test intended value, human learning, skills, and consequences over time. Reassess as conditions change. |

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

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0

# WHAT SHOULD WE LOOK AT?

Use these questions to find where to look first, then use the full assessment.

## For an organization or institution

Direction & accountability — Who owns this endeavor? Who can approve, change, pause, override, or stop it?
People & competence — Do users and supervisors have the learned judgment, practice, skills, and support 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 and harm? Are reporting channels usable, with named recipients, investigators, action owners, and validated corrections?
Security, safety & resilience — What happens when the AI, a vendor, an agent, an integration, or a critical dependency fails or behaves unexpectedly?
Enterprise integration — Has AI moved to team, cross-functional, or enterprise scale faster than stewardship?
Resources & economics — Who funds monitoring, maintenance, enhancement, human review, and retirement?
Human consequences — Who benefits? Who bears consequences? What happens to agency, capability, relationships, trust, opportunity, and human flourishing?

## For an individual

Individual responsibility does not remove provider, institutional, or public obligations. Children and users unable to recognize or manage risks need protection and support beyond their own verification skills.

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 judgment, accountability, agency, learning, independent thinking, and essential skills?
Do I know when to seek human expertise, escalate, refuse, correct, or stop?
Am I learning from outcomes? Correcting an answer or submitting feedback is not evidence that the system has been fixed.

## 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, 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 AI endeavor. Bring together its accountable leader, a knowledgeable user, technical expertise, and risk or stewardship responsibility. Check the preliminary gate and work through the eight steps. Select the three to five most consequential gaps, 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 three-domain stewardship portfolios and coverage-gap tools.

© 2026 by John Daniels and Prometheus Endeavor is licensed under CC BY 4.0