



RESPONSIBLE AI PRACTICE STANDARD

HUMAN-LED.
AI-ASSISTED.
ACCOUNTABLE BY DESIGN.

Our commitments to responsible, risk-based and human-led AI use.

Purpose

Possibility Engineering Consulting (PEC) believes artificial intelligence can materially expand the capacity of people and organizations to research, analyze, create, learn, decide, and execute. Increased technological capability must be accompanied by increased clarity regarding responsibility.

This standard establishes how PEC uses AI within our consulting practice and the principles applied when advising organizations on AI-enabled work. Our goal is not simply to use AI efficiently. It is to use AI deliberately, transparently, securely, critically, and accountably.

Our Core Position

Human-led. AI-assisted. Accountable by design. AI may contribute to research, synthesis, analysis, ideation, design, workflow development, and production. It does not replace professional accountability. PEC remains responsible for the interpretations, recommendations, methodologies, and final work it delivers to clients.

01 Human Accountability

A human professional remains accountable for PEC work.

AI systems are not accountable professional actors.

PEC does not delegate ultimate responsibility for professional judgment, client recommendations, organizational diagnoses, personnel decisions, strategic conclusions, ethical judgments, validation of material findings, or final deliverable approval.

Where AI makes a substantive contribution, a qualified professional reviews the work. Oversight includes the authority and practical AI systems are not accountable professional ability to question, correct, reject, or replace AI-generated content or recommendations.

02 Purposeful Use

PEC uses AI where it improves the work, not merely because it is available.

Before introducing it into a process, we consider the problem being solved, expected value, potential risk, appropriate alternatives, and the human capabilities that must remain involved.

Efficiency is one consideration, but not the only one. Quality, evidence, learning, confidentiality, relationships, and the effect on organizational capacity also matter. What can be automated is not automatically what should be automated.

03 Appropriate Human Oversight

Human oversight is proportionate to the sensitivity, uncertainty, complexity, and foreseeable consequence of the work. Routine production assistance generally requires less review than organizational diagnostics, confidential analysis, people-related assessments, financial implications, or consequential recommendations.

For consequential work, oversight may be needed throughout the process, not just at final sign-off. Where meaningful oversight cannot be provided, PEC will restrict, redesign, or decline the proposed use.

04 Verification and Evidentiary Integrity

AI output is not automatically treated as fact, evidence, or professional analysis.

PEC independently verifies material factual claims, calculations, source references, analytical findings, and interpretations where these support consequential conclusions or recommendations.

Our review may involve primary sources, original data, checked calculations, alternative explanations, assumptions, uncertainty, and methodological limitations. AI-generated citations must be verified. An AI system reviewing another AI system is not necessarily independent validation.

Unverified information is not knowingly represented as established fact. We distinguish evidence, interpretation, hypothesis, and AI-generated suggestion.

05 Confidentiality, Privacy, and Data Stewardship

Client confidentiality is a governing obligation irrespective of technology.

PEC evaluates information sensitivity, client permissions, contractual terms, data-use practices, retention, access and security protections before processing client information with AI.

Sensitive categories include personal and employee information, assessments, health-related or protected data, proprietary strategy, intellectual property, privileged material, and other confidential organizational information.

Appropriate controls include permitted tools and configurations, minimizing data shared, aggregation or de-identification when appropriate, and required client authorization. Permission to perform consulting work does not automatically authorize processing sensitive data through external AI providers.

Where adequate safeguards, contractual permissions, or lawful grounds are unavailable, PEC will not proceed with the proposed processing.

06 Transparency and Disclosure

PEC does not knowingly misrepresent materially AI-assisted professional work as exclusively human-produced where that impression would reasonably affect a client's understanding of the work. Disclosure is proportionate to the contribution, sensitivity, decision consequences, client agreement, and applicable requirements.

Incidental editing or formatting normally does not require individual disclosure. Substantive AI involvement may warrant engagement-level or deliverable-level disclosure. Processing sensitive client information may additionally require prior authorization; disclosure after the fact is not a substitute.

Our companion AI Use, Risk & Disclosure Framework explains the approach.

07 Human Judgment at Consequential Decisions

Automation does not automatically confer decision authority.

Where AI informs consequential decisions, PEC expects identifiable human accountability, a clear evidentiary basis, meaningful challenge and override, and attention to potential effects on individuals and organizations.

AI-generated classifications, rankings, scores, or recommendations are not independently sufficient grounds for consequential personnel decisions.

Avoiding the AI-in/AI-out loop. When AI helps produce work and another AI process evaluates it, apparent validation may mask shared errors or assumptions. PEC assesses whether independent evidence and meaningful human judgment remain present.

08 Bias, Fairness, and Organizational Context

AI systems can reflect or amplify limitations in data, measurement, models, prompts, and assumptions. PEC therefore evaluates their output in context, particularly when work concerns people, employment, leadership, assessment, performance, opportunity, and organizational power.

Appropriate review may examine representativeness, construct relevance, accessibility, potential disparate effects, alternative explanations, and relevant legal or professional obligations. A pattern detected by AI is not, by itself, an explanation of its cause.

09 Human Capability, not Dependency

PEC seeks to use AI to expand human and organizational capability rather than unnecessarily erode it. Some tasks that technology can perform are also experiences through which professionals develop judgment, expertise, relationships, and confidence.

When redesigning work, we consider what foundational practice people may lose, what new capabilities they need, and what learning opportunities should be intentionally retained or created, particularly for early-career professionals and emerging leaders.

This does not imply AI should always be withheld. It means capability development deserves deliberate design, alongside efficiency gains.

10 Responsible Experimentation and Continual Improvement

Responsible AI is an evolving discipline. PEC assesses tools and uses before consequential adoption, tests methods where appropriate, documents material AI involvement, addresses errors, maintains relevant fluency, and reviews important changes in capabilities, provider terms, and requirements.

When an application cannot be operated with sufficient reliability, oversight, confidentiality, or accountability, PEC will modify, suspend, or discontinue its use. Material incidents require appropriate assessment and corrective action.

11 Client Choice and Professional Boundaries

PEC respects client-specific AI limitations, contractual obligations, confidentiality expectations, and permitted uses. Where material AI use requires discussion or authorization, it will be addressed before the relevant work proceeds.

Client consent does not remove PEC's independent duty to exercise professional judgment or meet applicable professional and legal obligations.

12 AI In The Service of Possibility

AI adoption should not simply remove labor or accelerate existing processes. It should help organizations improve how work is performed, how people contribute, how decisions are made, how expertise develops, and what the organization is capable of achieving.

Decide. Design. Govern. Then Deploy. We believe responsible AI begins with organizational purpose, intentional work design, effective governance, and only then deployment.

Technology expands possibility. People remain responsible for what becomes of it.

Reference Frameworks

PEC draws on recognized approaches including the NIST AI Risk Management Framework, the NIST Generative AI Profile, and ISO/IEC 42001. These references inform practice; they do not signify PEC certification, accreditation, or complete conformity with a given standard. Applicable legal and client-specific requirements are assessed in context.

- NIST AI RMF: <https://www.nist.gov/itl/ai-risk-management-framework>
- NIST Generative AI Profile: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- ISO/IEC 42001: <https://www.iso.org/standard/42001>