PETRICHOR.

A CLIENT-READY AI ASSURANCE SYSTEM.

How high-trust firms prove responsible AI use without slowing responsible adoption

Executive brief
Prepared by: Petrichor
August 2026

CLARITY PRECEDES MOMENTUM.

1. THE QUESTION CLIENTS ARE STARTING TO ASK

> **How does AI touch our work, our information, and the decisions you make for us?**

For communications firms, law firms, and government contractors, that question is becoming commercially important.

Clients do not only want to know whether a firm has an AI policy. They increasingly want to understand:

- Where AI is and is not used in their work
- Whether sensitive information enters an AI system
- Which tools and vendors are approved
- Where human review and professional judgment remain mandatory
- How outputs are checked, sourced, and corrected
- What happens when a tool, workflow, or requirement changes
- Whether employees understand and follow the rules

Most firms can describe their intentions. Far fewer can produce a concise answer supported by current operating evidence.

That difference is the **AI assurance gap**.

Why a policy is not enough

A policy expresses intent. Assurance connects that intent to actual workflows, accountable people, approved tools, current evidence, and a repeatable review process.

Without that connection, client responses become slow, improvised, and difficult to defend.

2. THE OUTCOME

The objective is straightforward:

> **When a client or procurement team asks how AI touches sensitive work, the firm can respond within 24 hours with a clear position, documented controls, accountable owners, and evidence that the controls operate in practice.**

The visible result is a **client-facing AI Assurance Pack**. Behind it sits the operating system that keeps every statement accurate.

> **The pack is what the client can show. The operating system is what makes the pack true.**

WHEN A CLIENT ASKS…	THE FIRM CAN SHOW…
Where is AI used?	A bounded inventory of approved uses and workflows
What happens to our information?	Data classifications, tool restrictions, and data-flow evidence
Who remains responsible?	Named human owners, review requirements, and escalation points
How do you protect quality?	Workflow-specific checks, source requirements, and provenance rules
Are employees following the rules?	Role-specific activation and certification records
What happens when something changes?	Tool review, exception, incident, and evidence-refresh processes

This gives leadership a reliable external answer while creating clearer, safer pathways for useful AI adoption internally.

3. WHAT WORLD-CLASS ASSURANCE ADDS

The strongest AI management systems do more than document approved behavior. They maintain a living record of each important AI use, test it against defined thresholds, connect every external claim to evidence, and give clients meaningful control over how AI may affect their work.

BEST-PRACTICE MECHANISM	WHY IT MATTERS TO A CLIENT
Risk-tiered AI use register	The firm can distinguish prohibited, high-control, managed, and low-risk uses rather than treating all AI alike
Versioned use-case factsheets	Each important use has a purpose, owner, data boundary, provider/model record, limitations, tests, approvals, and next review date
Standards crosswalk	Controls can be related to NIST AI RMF, ISO/IEC 42001, and the CSA AI Controls Matrix without implying certification
Client rights and restrictions	A client can prohibit AI, limit it to approved uses or tools, require disclosure or consent, and establish matter-specific boundaries
Supplier shared-responsibility record	The firm can explain training use, retention, deletion, residency, subprocessors, logging, incident terms, and what responsibility remains with the firm
Task-specific evaluation	High-control uses are tested for the failures that matter in the actual workflow—not only against generic model benchmarks
Evidence freshness rules	An expired artifact or materially changed system can no longer support an external claim without review
Material-change and incident process	Model, terms, workflow, data, integration, or client-requirement changes trigger reassessment; incidents trigger containment, evidence preservation, and notification analysis
Independent challenge	At least annually, an independent reviewer tests the inventory, evidence, client claims, and operation of controls

A management system, not a binder

ISO/IEC 42001 uses a continual-improvement model: establish, operate, evaluate, correct, and improve. NIST's Generative AI Profile similarly emphasizes governance, content provenance, pre-deployment testing, and incident disclosure. The Cloud Security Alliance extends this into a detailed AI control and shared-responsibility model.

The practical implication is simple: an assurance pack must have an **as-of date, evidence status, material-change process, and next review date**. Otherwise it begins aging the moment it is issued.

4. WHAT THE ASSURANCE SYSTEM CONTAINS

COMPONENT	WHAT IT ESTABLISHES
1. Assurance boundary	An inventory of AI uses, tools, sensitive information, client obligations, teams, and workflows; precisely what the firm is and is not assuring
2. Risk-tiered use register	Prohibited, high-control, managed, and low-risk treatment based on information sensitivity, client impact, autonomy, reversibility, and judgment dependency
3. Versioned use-case factsheets	Purpose, owner, data, tool/provider/model, limitations, tests, client restrictions, approval status, change triggers, and next review date
4. Workflow controls	Permitted and prohibited uses, approved environments, human review, evidence requirements, and escalation triggers attached to actual work
5. Tool, data, and supplier governance	Training use, retention, deletion, residency, subprocessors, logging, incident terms, material changes, and shared responsibilities
6. Human accountability	Where AI may assist, where human judgment is mandatory, and who remains responsible for the final decision or output
7. Task-specific evaluation	Pre-use and post-change testing for grounding, accuracy, confidentiality leakage, instruction adherence, expert quality, unsafe inputs, and regression
8. Client pack and response library	A concise external position plus approved answers for questionnaires, RFPs, outside-counsel reviews, procurement, and contracts
9. Staff activation and certification	Role-specific scenarios, demonstrated understanding, recorded completion, and clear escalation routes
10. Continuous review and independent challenge	Management review, evidence refresh, corrective action, and periodic independent testing of the claims and controls

5. WHAT APPEARS IN THE CLIENT-FACING PACK

The pack is designed to be concise enough for an executive to send and specific enough for a client, general counsel, security team, or procurement reviewer to evaluate.

Core contents

1. Executive assurance letter with version, coverage date, and next review date
2. Responsible-AI position and explicit assurance boundary
3. Risk-tiered scope of AI use in client work
4. Sensitive-information handling statement
5. Approved-tool, training-use, retention, deletion, and vendor-oversight summary
6. Human-review and accountability model
7. Use-case evaluation and continuous-monitoring summary
8. Quality, evidence, and provenance controls
9. Client choice, restriction, consent, and disclosure pathway
10. Staff activation and certification statement
11. Exception, material-change, and incident-response summary
12. Standards crosswalk and independent-review status
13. Evidence index and freshness status
14. Explicit exclusions and contact pathway

Client-specific annexes

When needed, the core pack can be supplemented for:

- A particular client's AI or security requirements
- Outside-counsel guidelines
- Restricted information or tool environments
- RFP and procurement questions
- Contract-specific obligations
- A prime contractor's assurance requirements
- Client opt-out, approved-use, tool, data, retention, or notification requirements

The core position remains controlled. Annexes address legitimate differences without creating a new policy for every client.

6. THE EVIDENCE STANDARD

World-class assurance makes traceability visible. Every material external statement follows one evidence chain:

```
External claim
→ assurance boundary
→ applicable use case and risk tier
→ operating control
→ accountable owner
→ current evidence
→ test or review result
→ exception status
→ next review date
```

What counts as current evidence

An evidence artifact must identify:

- The claim and control it supports
- Its source and system of record
- Its accountable owner and reviewer
- The period and population it covers
- Its approval and expiry dates
- Known limitations or exclusions
- Open exceptions and corrective actions

If a provider changes its model, terms, retention behavior, subprocessor list, integration, or security position, affected evidence is reviewed before the external claim is repeated.

What a use-case factsheet shows

For each important use, a reviewer can see:

- Why the firm uses it
- Which client, matter, contract, or workflow it affects
- What data and tools it relies on
- What AI may produce and may never decide
- What failure modes were anticipated
- How performance and confidentiality were tested
- Which client restrictions apply
- Who approved it and who reviews the output
- What would suspend or reopen approval

This creates a practical audit trail without forcing a client to inspect the firm's entire internal environment.

7. HOW IMPLEMENTATION WORKS

STAGE	WHAT HAPPENS	PROOF GATE
1. Establish the boundary — typically four weeks	Interview leadership and workflow owners; inventory actual AI use and tools; review client demands; map sensitive information and data flows; risk-tier material uses; establish account, matter, or contract restrictions; baseline current performance	Leadership approves the scope, owners, workflows, risk tiers, restrictions, and external position before controls are built
2. Build the system — typically eight to ten weeks	Encode workflow controls; establish approved-use and data-handling rules; build factsheets; document provider responsibilities; run task-specific evaluations; define human escalation; assemble evidence; produce the pack and response library	Every external statement maps to an operating control, current evidence, a test or review result, and a named owner
3. Activate and prove — typically six to eight weeks	Run role-specific activation; record certification; conduct an executive tabletop; simulate a demanding client review; test a material change and incident-notification decision; remediate material gaps	The firm responds to a realistic assurance request within 24 hours without unsupported claims or critical evidence gaps
4. Continue assurance — ongoing	Maintain factsheets, evaluation results, vendor records, restrictions, evidence freshness, incidents, exceptions, and corrective actions; conduct management review and annual independent challenge	Material changes reopen affected approvals; leadership records corrective actions and residual-risk decisions

8. ONE SYSTEM, THREE SECTOR EDITIONS

The underlying architecture remains the same. Each edition changes the obligation library, workflow patterns, language, and challenge scenario.

	COMMUNICATIONS	LEGAL	GOVERNMENT CONTRACTING
Sensitive information	Embargoed announcements, crisis facts, executive vulnerabilities, strategy, media lists	Privileged communications, work product, matter files, PII, deal information	FCI/CUI boundaries, technical data, proposal data, program information
Typical review audience	Client general counsel, communications leadership, CISO, procurement	Client general counsel, outside-counsel management, InfoSec, privacy	Contracting organization, prime contractor, security/compliance, procurement
High-risk workflows	Crisis analysis, monitoring, briefing, messaging, media strategy, executive preparation	Research, drafting, discovery, due diligence, matter summarization, client advice	Capture, proposal, program research, technical writing, knowledge retrieval, reporting
Human boundary	Final strategy, factual approval, client counsel, public release	Legal judgment, advice, citation validation, privilege decisions, supervision	Contract interpretation, controlled-data decisions, representations, technical approval
Client control	Account-level disclosure, approval, tool, or no-AI restriction	Matter-level consent, opt-out, tool, data, billing, and outside-counsel restrictions	Contract-, system-, data-, prime-, and procurement-specific restrictions
Evaluation focus	Factual grounding, confidentiality, tone, source quality, and escalation	Citation validity, privilege/confidentiality, accuracy, supervision, and billing transparency	Controlled-data handling, source grounding, instruction adherence, and representation accuracy
Commercial outcome	Protect sensitive retainers and make responsible AI visible to clients	Retain institutional clients while protecting privilege, quality, and professional accountability	Support bid and contract readiness with a documented AI assurance layer

Important boundaries

- The communications edition does not guarantee factual accuracy or eliminate crisis risk.
- The legal edition does not provide legal advice or determine privilege.
- The government-contracting edition does not assess or certify CMMC, NIST, FedRAMP, or cybersecurity compliance.

The system complements qualified legal, privacy, cybersecurity, and compliance professionals. It does not replace them.

9. WHAT CHANGES FOR THE FIRM

BEFORE	AFTER
Leadership knows AI is being used but cannot describe the boundary	Approved uses, restrictions, owners, and evidence are visible
All AI uses are discussed as if they carry the same risk	Every material use is registered and assigned an assurance tier
Client answers are assembled from scratch	A controlled pack and response library support consistent answers
Policies sit apart from daily work	Rules are attached to specific workflows and decisions
Staff either avoid AI or improvise	Employees have clear approved pathways and escalation routes
Tool reviews happen inconsistently	New tools and material changes follow one review process
Vendor promises are treated as the control	Provider, platform, firm, and user responsibilities are documented separately
Human review is assumed	Accountability and review thresholds are explicit
Evidence is scattered or ages silently	Claims, controls, owners, tests, expiry dates, and exceptions are connected
Clients receive disclosure but little choice	Account-, matter-, or contract-level restrictions can be recorded and enforced
Governance is treated as a one-time project	Assurance is maintained, independently challenged, and improved as conditions change

What leadership gains

- A credible answer for demanding clients
- Visibility into actual AI use
- Clearer decisions about where AI should and should not be used
- Less repeated executive, legal, and security effort
- A foundation for responsible adoption without broad, unsupported promises
- A client-visible demonstration that confidentiality and expert judgment remain central
- Readiness for independent assessment or ISO/IEC 42001 certification when commercially justified

10. A USEFUL INTERNAL CONVERSATION

The following questions will quickly reveal whether an assurance gap exists:

1. Which client would be hardest to answer if it asked tomorrow how AI touches its work?
2. Where are employees already using AI with sensitive information?
3. Can leadership distinguish approved use, restricted use, and prohibited use by workflow?
4. What evidence could the firm produce today—not merely what policy could it describe?
5. Who approves new tools, client disclosures, exceptions, and incident responses?
6. What useful AI adoption is currently blocked because leadership cannot bound the risk?
7. Could the firm respond to a demanding client review within 24 hours?
8. Can a client opt out or impose a matter–, account–, or contract–specific restriction that staff will actually see and follow?
9. Which approved uses have current factsheets, test results, and next review dates?

A practical first step

Select one demanding client and three sensitive workflows. For each workflow, identify:

- What information enters the process
- Which tools may be used
- What AI is permitted to do
- What AI may never decide
- Who reviews the work
- What evidence proves the process was followed
- What client restriction, consent, or disclosure applies
- What change would force reassessment

If those answers are unclear or inconsistent, the firm does not yet have an AI assurance system. It has intentions.

ABOUT PETRICHOR

Petrichor helps high-trust organizations turn expert judgment into operating systems that are measurable, defensible, and usable in practice.

For AI assurance work, Petrichor connects client expectations, sensitive workflows, human accountability, staff activation, and operating evidence—so a firm's external position reflects how work is actually performed.

Petrichor
Clarity precedes momentum.

Reference foundations

The approach is informed by the NIST AI Risk Management Framework and Generative AI Profile, ISO/IEC 42001 AI management-system principles, the Cloud Security Alliance AI Controls Matrix, professional-responsibility guidance, current client and procurement expectations, and responsible-AI practices across communications, legal, and government-contracting environments.

The legal edition also reflects the Association of Corporate Counsel's current outside-counsel readiness signals: clear policies, named accountability, opt-in controls, systematic quality checks, and transparent billing.

Selected references

- NIST AI RMF: Generative AI Profile — https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
- ISO/IEC 42001 AI management systems — https://www.iso.org/standard/42001
- Cloud Security Alliance AI Controls Matrix v1.1 — https://cloudsecurityalliance.org/artifacts/ai-controls-matrix-v1-1
- ACC outside-counsel GenAI readiness questions — https://www.acc.com/resource-library/top-10-genai-transparency-readiness-questions-outside-counsel
- OWASP GenAI Incident Response Guide — https://genai.owasp.org/resource/genai-incident-response-guide-1-0/

This brief describes an operating and assurance approach. It is not legal advice, a cybersecurity assessment, or a compliance certification.