

# The Black Box Audit

KPMG Clara runs analytics across 100% of your journal entries. EY Astra drafts audit memo language from flagged conditions. Deloitte Omnia surfaces anomalies before the engagement team reviews them. Your audit committee approved an engagement letter that may not describe any of this — and the PCAOB has signaled that AI-assisted audit procedures carry the same documentation requirements as human-performed ones.

01

# Background

How AI became standard practice inside every Big Four audit — and what the governance architecture looks like from the client side.

## Four platforms. 100% journal entry coverage. One unresolved oversight question.

- KPMG Clara: AI audit platform deployed globally since 2017; significantly upgraded with GenAI capabilities in 2023–2025. Clara Analytics reviews 100% of journal entries (not just statistical samples); Clara Documents uses GenAI for contract and agreement extraction. Deployed across 80+ countries. Clara is the most mature enterprise-scale audit AI in public deployment.
- EY Astra: GenAI embedded into EY Canvas (the primary audit workflow platform) and EY Atlas (regulatory knowledge base). Astra drafts risk assessment narratives, synthesizes findings from structured data analysis, and surfaces anomalies for engagement team review. EY has publicly committed Astra to all audit engagements.
- Deloitte Omnia + DAN: Omnia is Deloitte's AI-powered audit platform running automated risk scanning across client financial data. Deloitte AI Navigator (DAN) overlays GenAI capabilities for memo drafting and regulatory research across 100,000+ professionals. Omnia flags patterns for engagement team review; DAN assists in documenting the team's response.
- PwC Halo + GenAI: Halo audit platform with GenAI capabilities for unstructured document processing — contracts, communications, board minutes. PwC also deployed Harvey AI (\$3B valuation) for deals and advisory work. Halo processes both structured financial data and unstructured client documents.
- PCAOB posture: The PCAOB issued Staff Guidance on AI in audit in 2024. AS 2301 (auditor's response to assessed risks) and AS 2315 (audit sampling) both have direct implications for AI-assisted procedures. The PCAOB has signaled that AI-assisted procedures require the same documentation standards as human-performed procedures — the tool does not change the documentation obligation.

## What AI actually does in your audit: three operational categories

- Journal entry analysis at full population scale: traditional statistical audit sampling reviews 2–5% of journal entries. Clara, Omnia, and Halo now run analytics on 100% of entries — anomaly detection, period-over-period pattern matching, Benford's Law analysis at scale. Higher population coverage changes the risk model but does not automatically increase assurance — assurance depends on how the engagement team evaluates and responds to what the model flags.
- Document extraction and synthesis: GenAI extracts key terms from contracts, leases, and agreements; identifies discrepancies between disclosed terms and financial statement treatment; drafts audit memo language from flagged conditions. The drafts are reviewed and approved by the engagement team — but the starting point is model output, and drafting speed changes how long engagement teams spend on memo development.
- Risk assessment support and anomaly triage: AI platforms cross-reference client-reported financials against industry benchmarks and prior-period results, flag statistical outliers, and generate ranked lists of items requiring engagement team attention. The team reviews the output and makes professional judgment calls — but the scope of what gets human attention is shaped by what the AI surfaces.
- What AI does not do: make the professional judgment calls that define the audit opinion. The auditor still assesses going concern, applies materiality, evaluates management estimates, and takes full professional responsibility for the opinion. AI changes how evidence is gathered, synthesized, and presented — it does not change who signs the report or who carries the liability.

## The client-side governance gap

- **Data access scope:** AI audit platforms require structured data export from the client's ERP, financial systems, and document repositories. The data flows to Big Four cloud infrastructure — primarily Azure (all four have announced Microsoft partnerships). Most engagement letters do not specifically identify the AI tools in use or describe the data handling obligations they entail beyond standard confidentiality provisions.
- **PCAOB documentation standard:** The PCAOB has signaled that when AI-assisted procedures make consequential scope decisions — what anomalies to flag, what patterns to surface — those decisions must be documented with the same specificity as human-performed procedures. "The AI reviewed the journal entries" is not audit documentation. What the AI procedure was, what parameters it used, what it found, and how the engagement team evaluated the findings must all be in the workpapers.
- **The dismissal problem:** AI journal entry platforms generate flags at a scale that human review queues cannot fully process at the same depth as targeted human investigation. The professional judgment decision about which flags to investigate and which to dismiss must be documented with a professional rationale — not just marked as "reviewed." "The AI did not flag it" is not a defensible answer to a PCAOB inspector asking why a pattern was not investigated.
- **Audit committee visibility gap:** Most audit committee agendas do not include a standing item for AI tool usage in the audit — what platforms were used, what procedures they supported, what data categories were processed, and how the engagement team reviewed AI output. This oversight gap is widening as AI scope inside audit engagements expands.

## Your audit engagement uses AI tools your audit committee has not explicitly reviewed. What oversight obligation follows — and have you met it?

Your Big Four engagement team is using AI tools for journal entry analysis, document extraction, and risk assessment synthesis. This is standard practice — but the engagement letter your audit committee approved may predate these tools or describe them in general language that does not reflect current operational scope or data access.

The PCAOB has established that AI-assisted audit procedures carry the same documentation requirements as human-performed procedures. Your audit committee's oversight responsibility includes understanding how AI is used in your audit, what data flows to auditor infrastructure, and whether the documentation standard the PCAOB expects is being met.

This is not a question about whether your auditor is competent or trustworthy. It is a question about whether the governance framework your audit committee operates under was designed for the audit your organization is actually receiving — or for the audit that existed before AI became standard practice.

## Four governance postures.

### Option A

#### **Continue current posture — engagement letter and audit committee agenda unchanged**

Defensible if your engagement letter explicitly addresses AI tool usage, data handling, and PCAOB documentation standards. Not defensible if AI scope has expanded beyond what the current letter describes — and for most pre-2023 engagement letters, it has.

### Option B

#### **Recommended**

#### **Request AI tool disclosure addendum to engagement letter — annual update as tools change**

Require the engagement team to identify AI platforms in use, procedures they support, data categories processed, and data retention and deletion terms. Does not limit auditor flexibility. Creates the audit committee visibility the oversight function requires.

### Option C

#### **Require AI procedure documentation in workpapers to PCAOB standard**

Request that AI-assisted procedures be documented with the same specificity as human-performed procedures — tool, parameters, output, professional judgment applied. Appropriate for PCAOB-regulated issuers, companies with prior PCAOB findings, and any issuer where inspection risk is elevated.

### Option D

#### **Commission independent AI audit review**

Have internal audit or an independent third party assess whether the Big Four AI tools in use meet PCAOB documentation standards. Highest governance posture. Right for issuers with significant deficiency history or where the audit committee cannot independently evaluate AI disclosure from the engagement team.

## **Request the AI disclosure addendum. Add AI oversight to the audit committee agenda. Ask one specific question.**

Add a standing item to audit committee meeting agendas: what AI tools did the engagement team use in this period, what procedures did they support, what data categories were transmitted to auditor infrastructure, and what was the engagement team's review protocol for AI output? This is a 10-minute briefing item that closes the structural oversight gap.

Request an AI tool disclosure addendum to the engagement letter specifying: names of AI platforms used in the engagement, specific procedures each platform supports, data categories transmitted to auditor cloud infrastructure, and data retention and deletion terms at engagement end. Require annual update as tools change.

Ask the engagement partner one specific question: if PCAOB inspects this audit and asks why a flagged anomaly was not investigated, what documentation exists recording the professional judgment rationale — and where in the workpapers is it? The answer tells you whether AI-assisted procedures are being documented to the standard the PCAOB has signaled it expects.

Audit your data handling terms. Engagement letters written before AI audit tools were deployed at scale — pre-2023 is a strong signal — may not accurately describe current data access and processing. The data categories flowing to auditor infrastructure today are materially broader than what was in scope when most current engagement letters were drafted.

For your next audit RFP, include AI governance as an explicit evaluation criterion: what platforms are used, how AI output is reviewed and documented, what the engagement team's override protocol is when the team disagrees with an AI flag, and how AI-assisted procedures are documented in workpapers to AS 2301 and AS 2315 standards.

## Five material risks.

1.

### **PCAOB inspection exposure from undocumented AI-assisted procedures**

The PCAOB has signaled that AI-assisted audit procedures require documentation to the same standard as human-performed procedures. An audit where AI tools made consequential scope decisions — what anomalies to flag, which flags to dismiss — without documented professional judgment rationale creates inspection exposure that did not exist in purely human-performed audits. The documentation obligation does not diminish because the procedure was faster.

2.

### **Data access scope expanding beyond engagement letter terms**

AI audit platforms have progressively expanded the categories of client data they access: structured financial data (established), contract and agreement terms (2022–2023), unstructured communications and board minutes (2024–2025). Engagement letters drafted before each expansion may not accurately describe current practice. The data your auditor's AI platform processes today may exceed what your audit committee approved.

3.

### **The black box dismissal — AI flags not investigated and not documented**

AI journal entry platforms generate flags at a volume that engagement teams triage rather than fully investigate at equal depth. The professional judgment decision about which flags to pursue and which to dismiss must be documented with a professional rationale. A workpaper that records AI flag volume but not the dismissal rationale for each dismissed flag is not compliant with the PCAOB's documentation standard for AI-assisted procedures.

4.

### **Foundation model concentration across all four Big Four firms**

KPMG, EY, Deloitte, and PwC have all announced Microsoft Azure partnerships for their AI audit platforms. If all four run on Azure OpenAI infrastructure, a correlated failure — platform outage, regulatory restriction on AI in audit, model update that changes anomaly detection behavior — affects all four audit firms simultaneously. Enterprise issuers whose previous and current auditors both rely on the same infrastructure have not assessed this concentration risk.

5.

### **Engagement team AI capability variance — audit quality risk not yet in PCAOB inspection findings**

AI audit tools are deployed firmwide, but engagement team capability to evaluate,

## Six questions for your audit committee and CFO.

1. Does your current engagement letter identify the AI platforms in use in your audit, the procedures they support, and the data categories transmitted to your auditor's cloud infrastructure — or does it use general confidentiality language that predates AI audit tools?
2. Has your audit committee received a briefing from the engagement partner on AI tool usage in your audit in the past 12 months — and is there a standing agenda item to receive that briefing annually as tools change?
3. If PCAOB inspects your last audit, can the engagement team produce documentation explaining the professional judgment rationale for each AI-flagged anomaly that was dismissed without full investigation — and is that rationale in the workpapers?
4. When did your organization last review the data handling terms in the engagement letter — and do they specifically address AI-related data access, cloud infrastructure, and data deletion at engagement end?
5. Did your most recent audit RFP include AI governance as an evaluation criterion — and do you know which AI platforms each firm uses, what procedures they support, and how they document AI-assisted procedures?
6. Who on your audit committee has sufficient understanding of AI audit platforms to evaluate an AI tool disclosure from the engagement team — and if the honest answer is no one, is that a capability gap the committee should address before the next renewal?

## Big Four AI Audit Platform Comparison

- KPMG Clara: most mature enterprise audit AI — 100% journal entry population analytics, GenAI document extraction, 80+ country deployment. Most transparent about platform capabilities in public materials. Clara's anomaly detection scope is well-documented; Clara's GenAI memo drafting capabilities are less so. Best starting point for RFP AI governance questions.
- EY Astra + Canvas: deepest workflow integration — Astra is embedded in Canvas, the primary engagement workflow, not a separate overlay. This means AI-assisted procedures are harder to separate from human procedures in workpaper documentation, and harder to scope in an engagement letter addendum. Request explicit procedure-level disclosure.
- Deloitte Omnia + DAN: broadest enterprise AI footprint — DAN reaches 100,000+ professionals across advisory and audit, making it difficult to scope which capabilities are active in a specific audit engagement. Omnia's structured data analytics are well-documented; DAN's role in audit memo drafting is less publicly defined. Require engagement-specific scoping.
- PwC Halo + Harvey: two-platform architecture means data scope and documentation obligations differ by procedure type. Harvey's legal AI capabilities are primarily active in deals and advisory — audit-specific GenAI capabilities are less publicly documented. Confirm which tools are active in your engagement specifically.
- Key RFP question: ask each firm to describe, in writing, which AI tools are active in your specific engagement, what procedures they support, and how those procedures are documented to PCAOB AS 2301 and AS 2315 standards. The specificity of the answer is itself an evaluation criterion.

AI INSIGHT LAB

# The Black Box Audit

AI Insight Lab publishes The Deployment Memo for enterprise executives who have to decide, not just read. One deployment dissected every Tuesday.

Subscribe at [aiinsightlab.cloud](https://aiinsightlab.cloud)