CherryHill
ADVISORY

AI Part 3:

The Auditor's Edge: Using AI Without Losing What Makes You Irreplaceable

Part 3 of 4 in the AI Certificate Program

April 15 | 12 PM ET

1 CPE Credit

www.cherryhilladvisory.com

Mike Levy

CIA, CRMA, CISA, CISSP, CDPSE, MBA

CEO & Managing Principal, *Cherry Hill Advisory*

Connect with Mike on LinkedIn

- 20+ years of experience in Internal Audit, Cybersecurity, and Risk Management
- Former Chief Audit Executive and Deloitte Consultant
- IIA International Internal Audit Standards Board Member
- Past Chair (2023-2024), IIA Northern American Board; Global Board Director
- Expert in aligning audit functions with evolving cyber risk landscapes

Midship

Kieran Taylor

MEng Computer Science, Cornell University · Y Combinator S24

Co-Founder & CEO *Midship*

Connect with Kieran on LinkedIn

CEO and Co-founder of Midship, AI agents that fully automate SOX testing and produce audit-ready work papers

Former senior engineer at Instacart (IPO prep) and Amazon

Lead engineer on AI product acquired by HubSpot

Y Combinator S24 alumnus; raised $4.15M seed round led by Costanoa Ventures

Deep expertise in audit automation, data infrastructure, and AI-powered financial controls

CherryHill
ADVISORY

Question and Answer (Q&A)

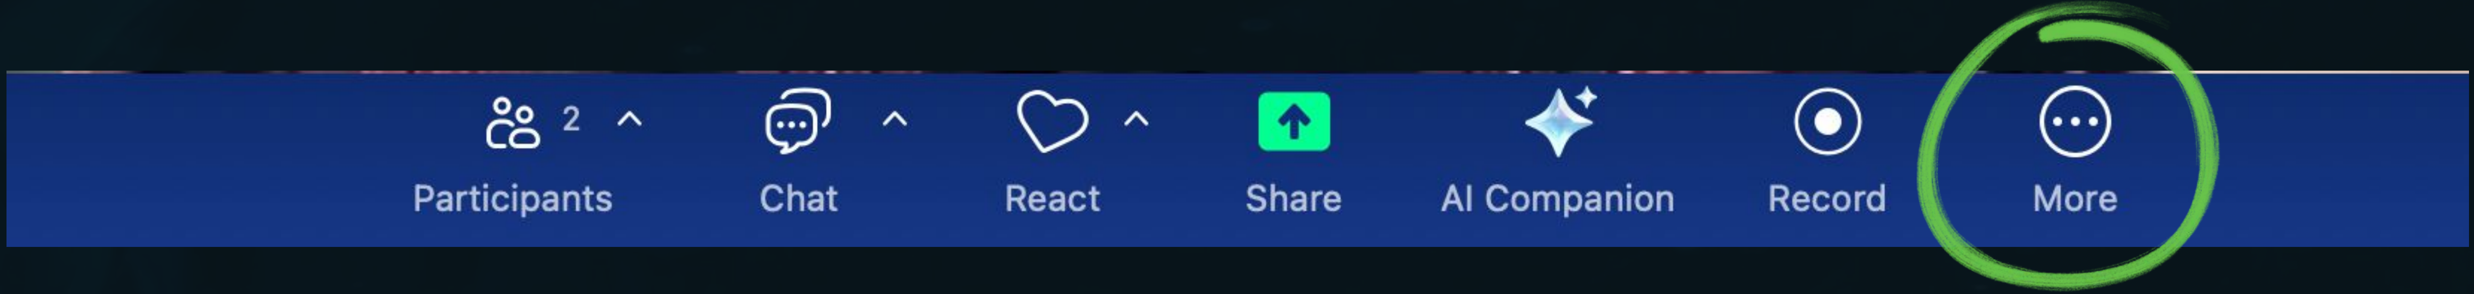

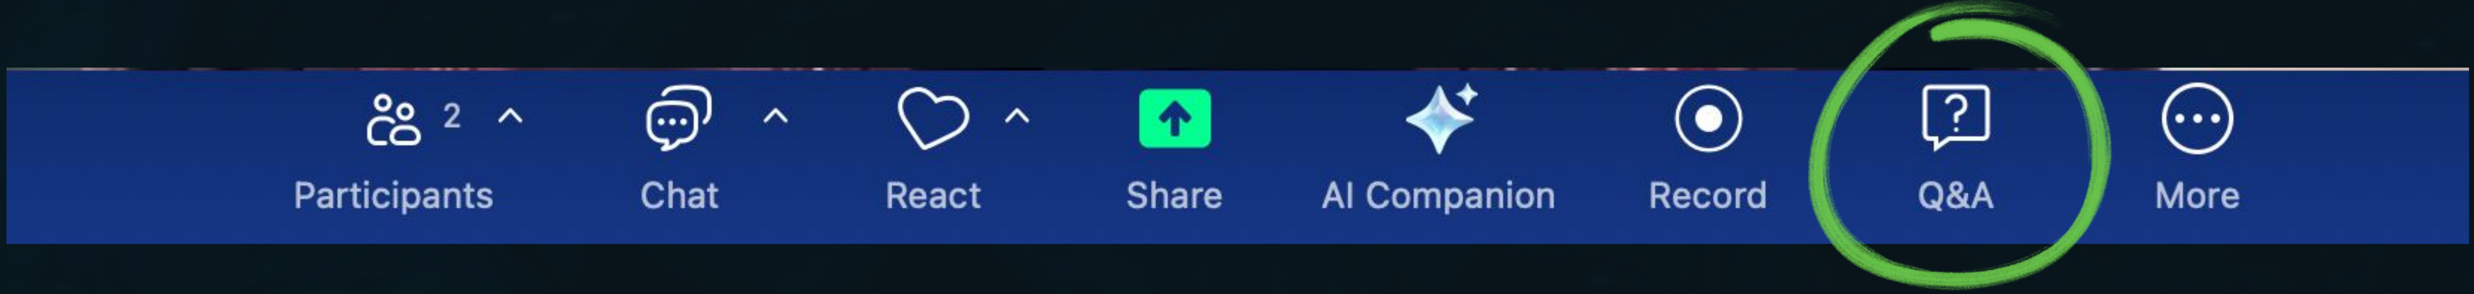

For your questions, the Q&A icon might show up in your screen. If not, you may find it by clicking the three dots with "More."

You have the option to ask anonymously.

What You Will Be Able to Do After This Session

01

Apply a structured scoping approach to define the audit population, identify AI touchpoints, and document key inputs, model characteristics, and decision outputs for an AI-enabled process.

02

Evaluate which testing procedures are appropriate for AI-enabled processes, including data quality tests, model governance checks, output sampling, and human-override control validation.

03

Differentiate AI-specific risk categories from traditional IT risks and connect them to core control objectives for accuracy, completeness, authorization, and validity in an audit workpaper.

Field of Study: Auditing | 1.0 CPE Credit | Intermediate Level | Prerequisite: Parts 1-3

Resources for this Session

CHERRY HILL ADVISORY'S AI AUDITING GUIDE

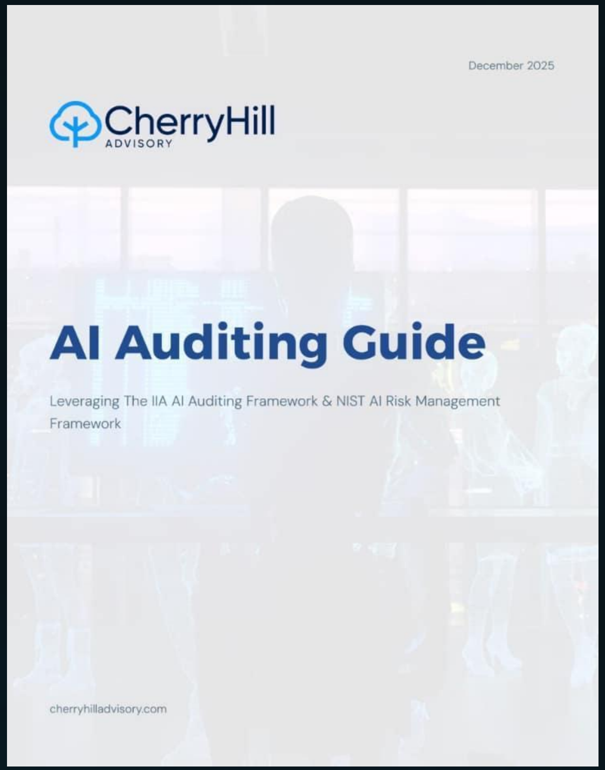

THE IIA'S GLOBAL INTERNAL AUDIT STANDARDS

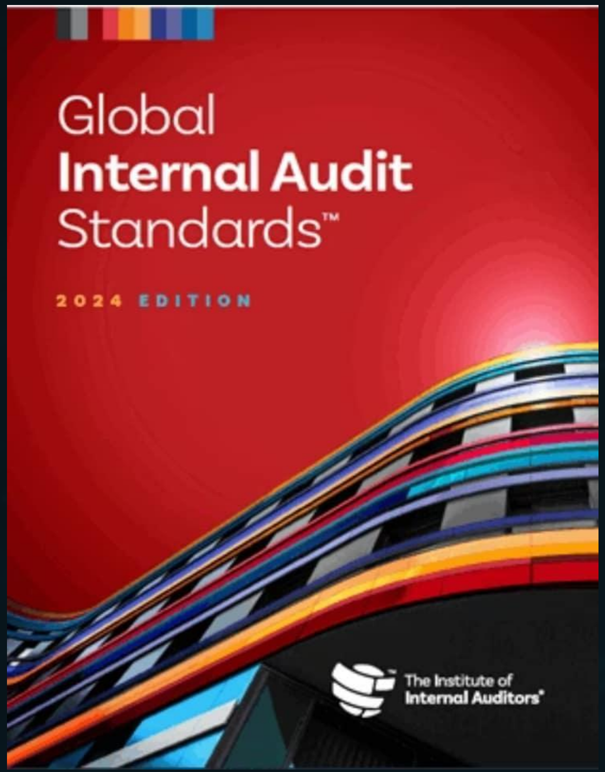

NATIONAL REGISTRY OF CPE SPONSORS

CPE Credit Information

- NASBA CPE Requirement: 50 minutes of attendance per hour, plus 3 poll questions.
- Certificates will be distributed mid next week or earlier.

For any questions or concerns, please email cpe@cherryhilladvisory.com.

Today's Agenda

1. **Context**
   Where the profession actually is, and what auditors are doing right now that they may not have fully thought through

2. **AI in Practice**
   The specific places in planning, testing, and reporting where AI earns its place, with step-by-step examples you can use Monday morning

3. **Guardrails**
   What you cannot hand off to a tool, and why the standards already answer this

4. **Documentation**
   The five-element standard that makes AI work defensible, including how to document the logic AI followed and the review steps you took

5. **Prompts + Workpaper Language**
   Ready-to-use review prompts for building an audit trail, plus sample workpaper language for the three most common AI-assisted scenarios

The Exposure Stat

41%

Using Generative AI

of internal audit functions report using generative AI in audit activities

59%

Not Yet Using It

are still not using AI but their organizations almost certainly are

6%

Using AI in Fieldwork

of functions use AI often or extensively in fieldwork the phase where it matters most

Source: Internal Audit Foundation, 2025 North American Pulse of Internal Audit, March 2025.

POLL 1 OF 4

When you use an AI tool to draft or analyze something in an audit engagement, how often do you go back to the underlying evidence to verify the output before including it in your workpaper?

- **Every time.** I treat AI output the same way I treat a third-party assertion.
- **Most of the time.** I verify the critical items but not every output.
- **Sometimes.** It depends on how confident I am in the tool.
- **Rarely.** If the output looks reasonable, I use it.
- **I have not thought about this** as a verification step.

The Gap Is Where the Risk Lives

Adoption Is Outpacing Governance

Many functions using AI right now have no written policy, no documentation standard, and no process for assessing independence implications.

That is not a criticism. It is a description of where the profession is.

The Window Is Closing

Adoption is outpacing governance across the profession. The gap between what teams are doing with AI and what their policies permit, or even acknowledge, is widening every quarter.

The window to get ahead of it is closing.

You Already Know How to Do This

The Methodology Has Not Changed

Auditors have always defined the criteria, evaluated the evidence, and owned the conclusion. AI handles the volume in between.

That is the entire framework and you have been doing the hard part for your whole career.

The Tool Has Changed

AI does not rewrite the audit process. It accelerates the parts that consume time without requiring judgment including data sorting, pattern detection, first-draft generation.

The hard part, the judgment, has always been yours. It still is.

AI in Audit: Current Applications

AI is rapidly transforming the audit landscape, moving beyond theoretical discussions to practical applications across various stages of the audit process. Here's how it's being deployed today:

Integrated Third-Party Software

Off-the-shelf audit tools now embed AI for capabilities like anomaly detection, predictive analytics, and automated testing procedures.

Custom-Built Applications

Organizations are developing tailored AI solutions to address specific audit needs, leveraging their unique data for specialized analysis and insights.

Generative AI for Content

Auditors utilize LLMs to draft initial reports, summarize complex documentation, and generate preliminary risk assessments, speeding up content creation.

AI-Enhanced RPA and Analytics

AI significantly improves Robotic Process Automation and Analytics, enabling more intelligent automation of repetitive tasks like data extraction, reconciliation, and control validation.

AI in Audit Planning: Where It Earns Its Place

Manual Approach	AI-Assisted Approach
Manually review each entity, cross-reference prior audits, and update ratings one by one	Export universe; AI flags rating changes with rationale; auditor reviews and overrides
Sort and filter findings tracker; identify patterns by hand across multiple cycles	Upload findings tracker; AI identifies aging patterns, departmental concentration, and recurring root causes
Pull reports, build pivot tables, and scan for outliers manually	Provide data extracts; AI identifies outliers, trends, and anomalies for audit attention
Read prior reports and regulatory updates; draft briefing notes from scratch	Feed AI prior report, management responses, and regulatory updates; receive a structured briefing summary

Source: CHA AI Auditing Guide, December 2025; GIAS Standards 9.1–9.4

CherryHill ADVISORY

AI in Fieldwork: The Testing Phase

Manual Approach	AI-Assisted Approach
Export data and manually stratify by risk criteria — amount, frequency, timing, user	Export full transaction population; AI stratifies by risk criteria and outputs a targeted sample design
Define test attributes and manually check each sample item against source documents	Give AI the sample data and source documents; AI performs initial attribute checks and flags exceptions for auditor validation
Review exceptions individually; manually group by approver, department, or time period	AI clusters exceptions by approver, department, time period, and access level; auditor determines root cause
Manually sort, label, and match unstructured evidence files to test samples	Upload unstructured files; AI parses, classifies, and matches documents to test samples in minutes

Source: CHA AI Auditing Guide, December 2025; GIAS Standards 11.1–11.4

CherryHill ADVISORY

POLL 2 OF 4

In your last completed audit engagement, how did you primarily use AI (if at all)?

- **Draft generation.** I used AI to generate a first draft of something (finding, memo, summary) that I then edited.
- **Data processing.** I used AI to process or analyze data (population analysis, exception patterns, trend identification).
- **Research and preparation.** I used AI to research or prepare (background research, standard interpretation, interview prep).
- **Something else** not listed here.
- **I did not use AI** in my last completed engagement.

AI in Reporting: From Draft to Delivery

Manual Approach	AI-Assisted Approach
Write each finding from scratch using condition, criteria, cause, and effect	Provide the four elements; AI drafts the finding; auditor reviews for accuracy, tone, and organizational sensitivity
Read the full audit report and manually write a condensed summary for audit committee review	Upload the full report; AI generates a two-paragraph board-level summary; auditor confirms it reflects professional opinion
Draft management response templates individually, assigning owners and timelines by hand	AI generates a suggested response template addressing the recommendation, owner, and proposed timeline
Review each completed report and manually compile board-level summaries and finding status updates	Feed reports and findings tracker to AI; receive two-sentence summaries, highest-risk findings, and response status for each report; CAE applies organizational context

Source: CHA AI Auditing Guide, December 2025; GIAS Standards 13.1–13.7

Three Guardrails

01 Data and Confidentiality

What goes into the tool is governed by policy. Before submitting: Is the tool approved by IT and legal? What is the data classification? Does the agreement prohibit training? If you cannot answer all three, do not submit.

02 Independence and Objectivity

Accepting AI output without independent validation creates a potential impairment. The test: Can you explain the conclusion without referencing the AI output?

03 Conclusion and Judgment

AI can draft, sort, calculate, summarize. AI cannot conclude. Every finding, opinion, and risk rating must be the auditor's determination based on evidence.

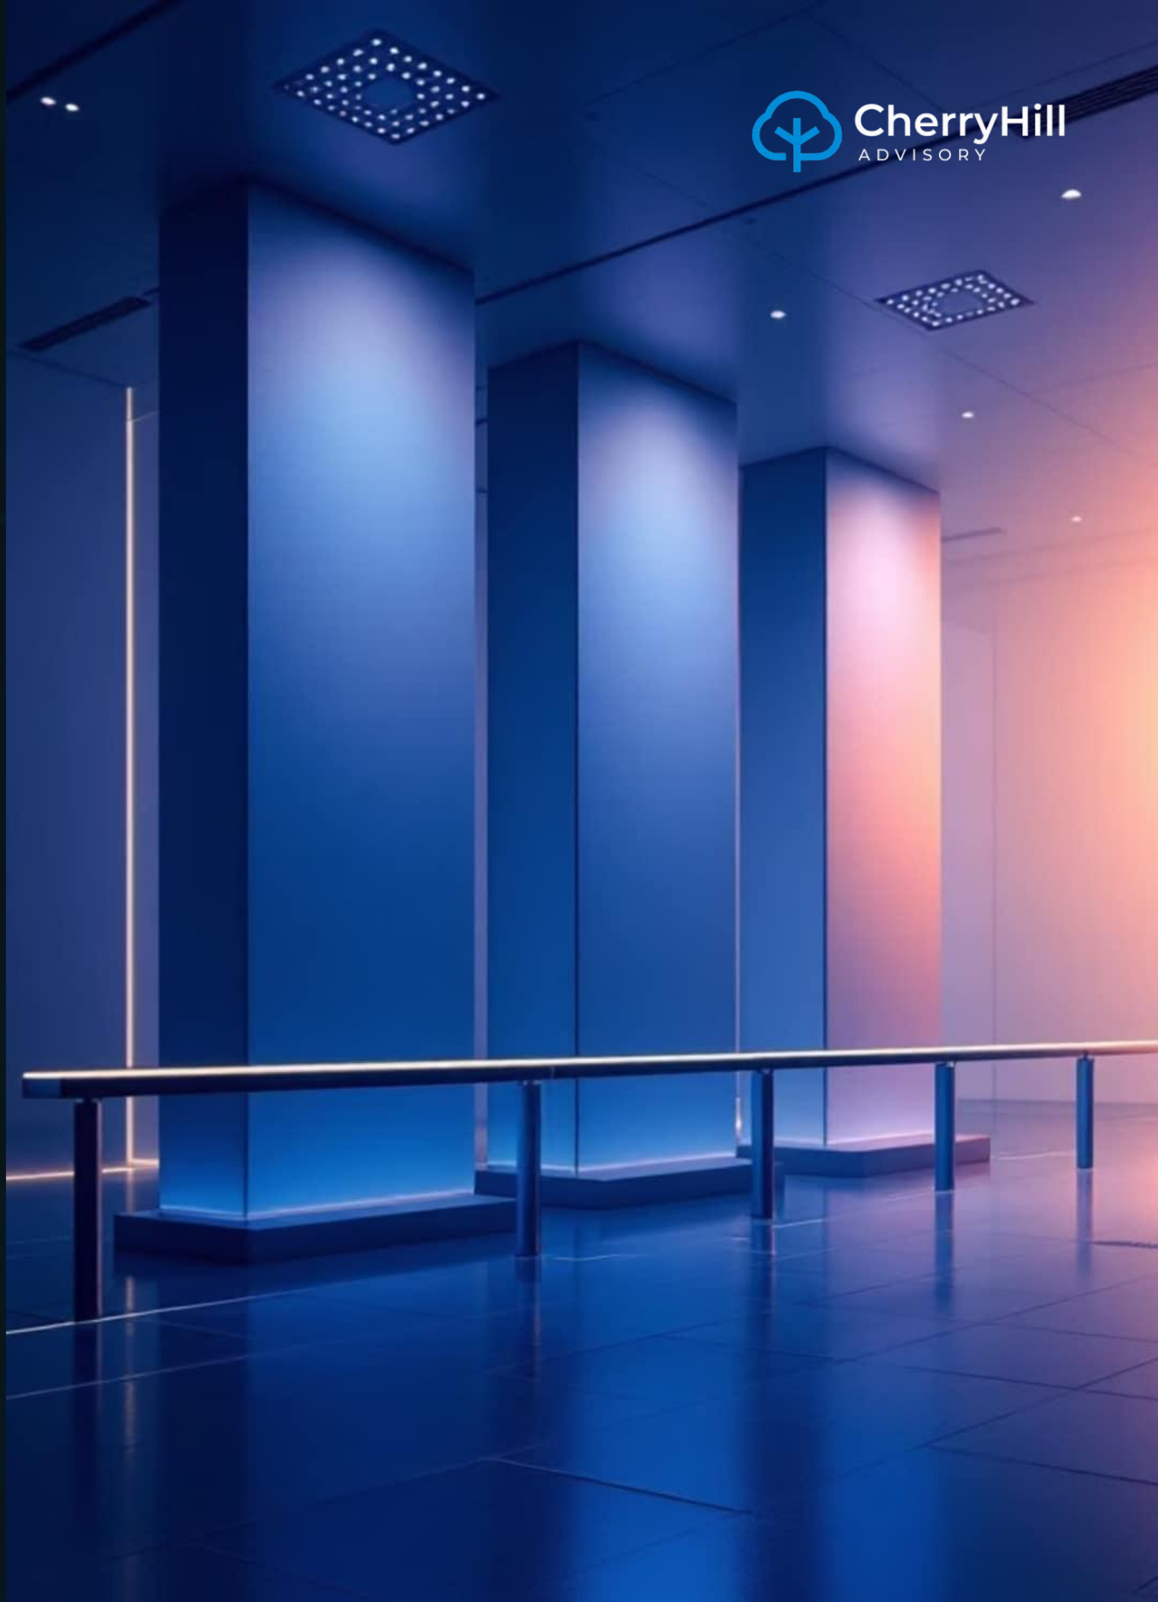

Data Confidentiality: What Goes In Is Governed by Policy

Authorization

Has your organization's IT and legal function approved this specific tool for use with organizational data?

Data Classification

What is the sensitivity level of the content you are about to submit? Client-identifiable information, audit evidence, control descriptions, and financial data each carry classification requirements.

Training Prohibition

Does your enterprise agreement explicitly prohibit the tool from using submitted content to train its model? If unknown, treat the tool as a public disclosure channel.

Retention and Storage

Where is the data stored? For how long? Who else at the provider can access it?

> The standard is not whether you trust the tool. The standard is whether you can demonstrate compliance with your organization's data governance policy.

POLL 3 OF 4

Have you ever submitted client-identifiable information, audit evidence, or confidential organizational data into a commercial AI tool (ChatGPT, Copilot, Gemini, Claude, etc.) as part of your audit work?

- **Yes, and my organization had approved** the tool and use case.
- **Yes, but I was not sure** whether the tool was approved.
- **Yes, and I now realize** it may not have been appropriate.
- **No, I have been careful** to avoid submitting confidential data.
- **I am not sure** whether what I submitted would be classified as confidential.

The Five-Element Documentation Standard

Tool Identification

Which AI tool? Which version? What access level (free, enterprise, API)? Different tiers have different data handling policies.

Prompt Documentation

What did you ask the tool to do? Document the actual prompt. If you refined across iterations, document the final version and note iteration occurred.

Output Retention

What did the tool return? Retain a copy; screenshot, export, or paste. The output is evidence of what the tool produced. It is not evidence of what is true.

Validation Description

What did you do to verify? Trace to source data? Test a sample against evidence? Compare to independent calculation? Document specific steps.

Conclusion Ownership

Explicit statement: "The conclusions in this workpaper are based on the auditor's independent evaluation of the evidence. AI-generated content was validated against [specific sources] before inclusion."

Source: GIAS Standards 11.1, 14.2; CHA AI Auditing Guide

AI-ASSISTED RISK ASSESSMENT

Key Prompting Principles for Auditors

These principles provide a framework for effectively leveraging AI in audit work, ensuring both efficiency and rigor.

Be specific about your role and context.	Define a structured output format for audit-ready results.	Always challenge the output.
Never skip human judgment. Document your prompts.	Specifies industry and organization context.	Explicitly requests emerging vectors that generic prompts miss.

ANATOMY OF A PROMPT TO SUPPORT FRAUD ANALYSIS

Dissecting the Prompt: Key Components

Understanding the structure of an effective AI prompt is crucial for extracting valuable, actionable insights, especially in complex areas like fraud risk.

Define AI's Persona

Instruct the AI on the role and expertise it should embody (e.g., "experienced internal audit and risk management expert with deep expertise in fraud risk"). This shapes its perspective and response style.

Specify Task & Context

Clearly articulate the core task (e.g., "build a comprehensive fraud risk universe") and the organizational context (e.g., "for a mid-size regional bank"), including specific scope parameters.

Detail Output Requirements

Prescribe the exact format and content for each output element (e.g., "for each scenario, provide: likelihood, financial impact, control gaps, and audit procedures").

Set Key Focus Areas

Direct the AI to prioritize specific considerations or constraints (e.g., "focus particularly on emerging digital fraud vectors... including AI-enabled fraud") to ensure relevance and depth.

ANATOMY OF A PROMPT TO SUPPORT FRAUD ANALYSIS

Prompt Review Discussion

```
"You are an experienced internal audit and risk management expert with deep expertise in
fraud risk. I need you to help me build a comprehensive fraud risk universe for a mid-size
regional bank. Please identify the top 15 to 20 fraud risk scenarios, organized by risk
category. For each scenario, provide: (1) likelihood — high, medium, or low; (2) potential
financial impact; (3) key control gaps that typically exist; and (4) two to three
recommended audit procedures. Focus particularly on emerging digital fraud vectors that may
not appear in traditional fraud risk assessments, including AI-enabled fraud."
```

POLL 4 OF 4

After seeing the five-element documentation standard today, which element is the one you are most likely to have been missing from your AI-assisted workpapers?

- **Tool identification**: which tool, which version, which access level
- **Prompt documentation**: what I asked the tool to do
- **Output retention**: a record of what the tool returned
- **Validation description**: what I did to verify the output
- **Conclusion ownership**: an explicit statement that the auditor, not the AI, owns the conclusion

CherryHill ADVISORY

Three Takeaways

01 AI Is the Tool. Judgment Is Your Edge.

AI accelerates volume. Risk universe refreshes, exception analysis, finding drafts, executive summaries. It does not replace skepticism, independence, or accountability. The auditor's value has always been in the judgment. AI takes the legwork. You keep what matters.

02 Prompt Well. Review Harder. Document Both.

An effective prompt gives AI a role, a task, constraints, and an output format. But the prompt is only half the work. What you do after, challenging the output, validating against evidence, documenting the logic, is what makes the work defensible. Both steps belong in the workpaper.

03 Five Elements. Every Time.

Tool and version. Prompt used. Output retained. Validation steps taken. Conclusion owned by the auditor. If those five are in the workpaper, you can defend anything in it. If any one is missing, you have a gap, and so does your supervisor, your QA reviewer, and your regulator.

What's Actually Happening at the Edge: Fireside Chat

What Breaks First

What breaks when audit teams start using AI in actual engagements, and what Midship has learned watching it happen at scale.

Practical Tips for Monday Morning

What a complete AI-generated workpaper contains and how to evaluate any AI tool for audit use, things you can act on immediately.

Where the AI Landscape Is Heading

From chatbot prompting to agentic AI that follows audit plans autonomously, what the next phase of AI in internal audit actually looks like.

CherryHill
ADVISORY

LEARNING NEVER STOPS

UPCOMING CPE WEBINARS, REGISTER HERE

AI PART 3:

Using AI Without Losing What Makes You Irreplaceable

Resources

CHERRY HILL ADVISORY'S AI AUDITING GUIDE

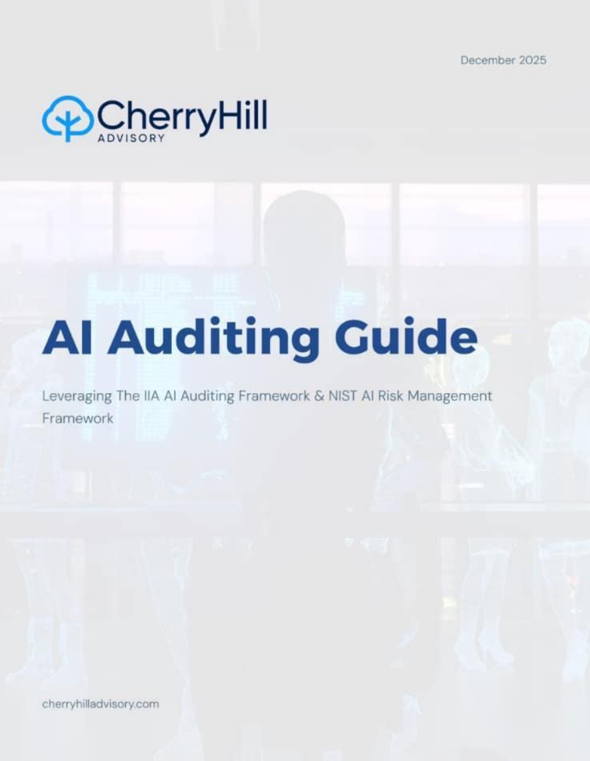

THE IIA'S GLOBAL INTERNAL AUDIT STANDARDS

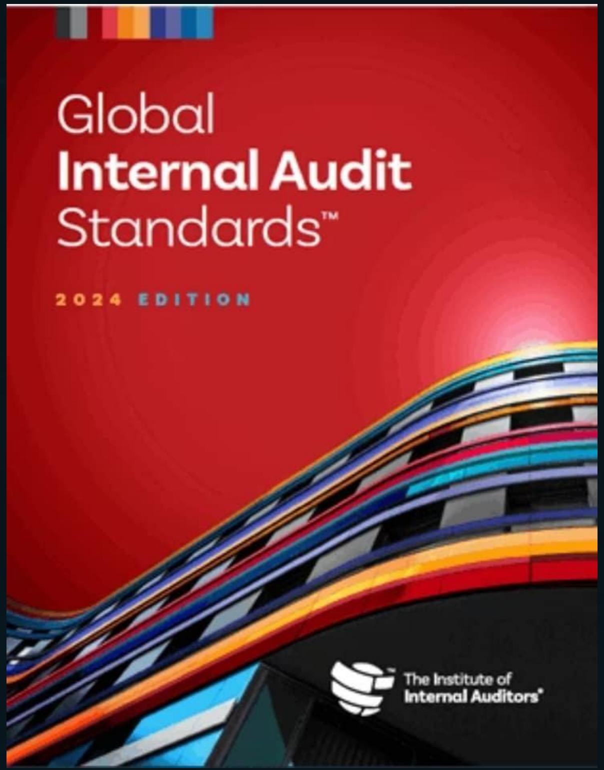

CherryHill ADVISORY

Thank you!

Connect with us on LinkedIn

MIKE LEVY

KIERAN TAYLOR