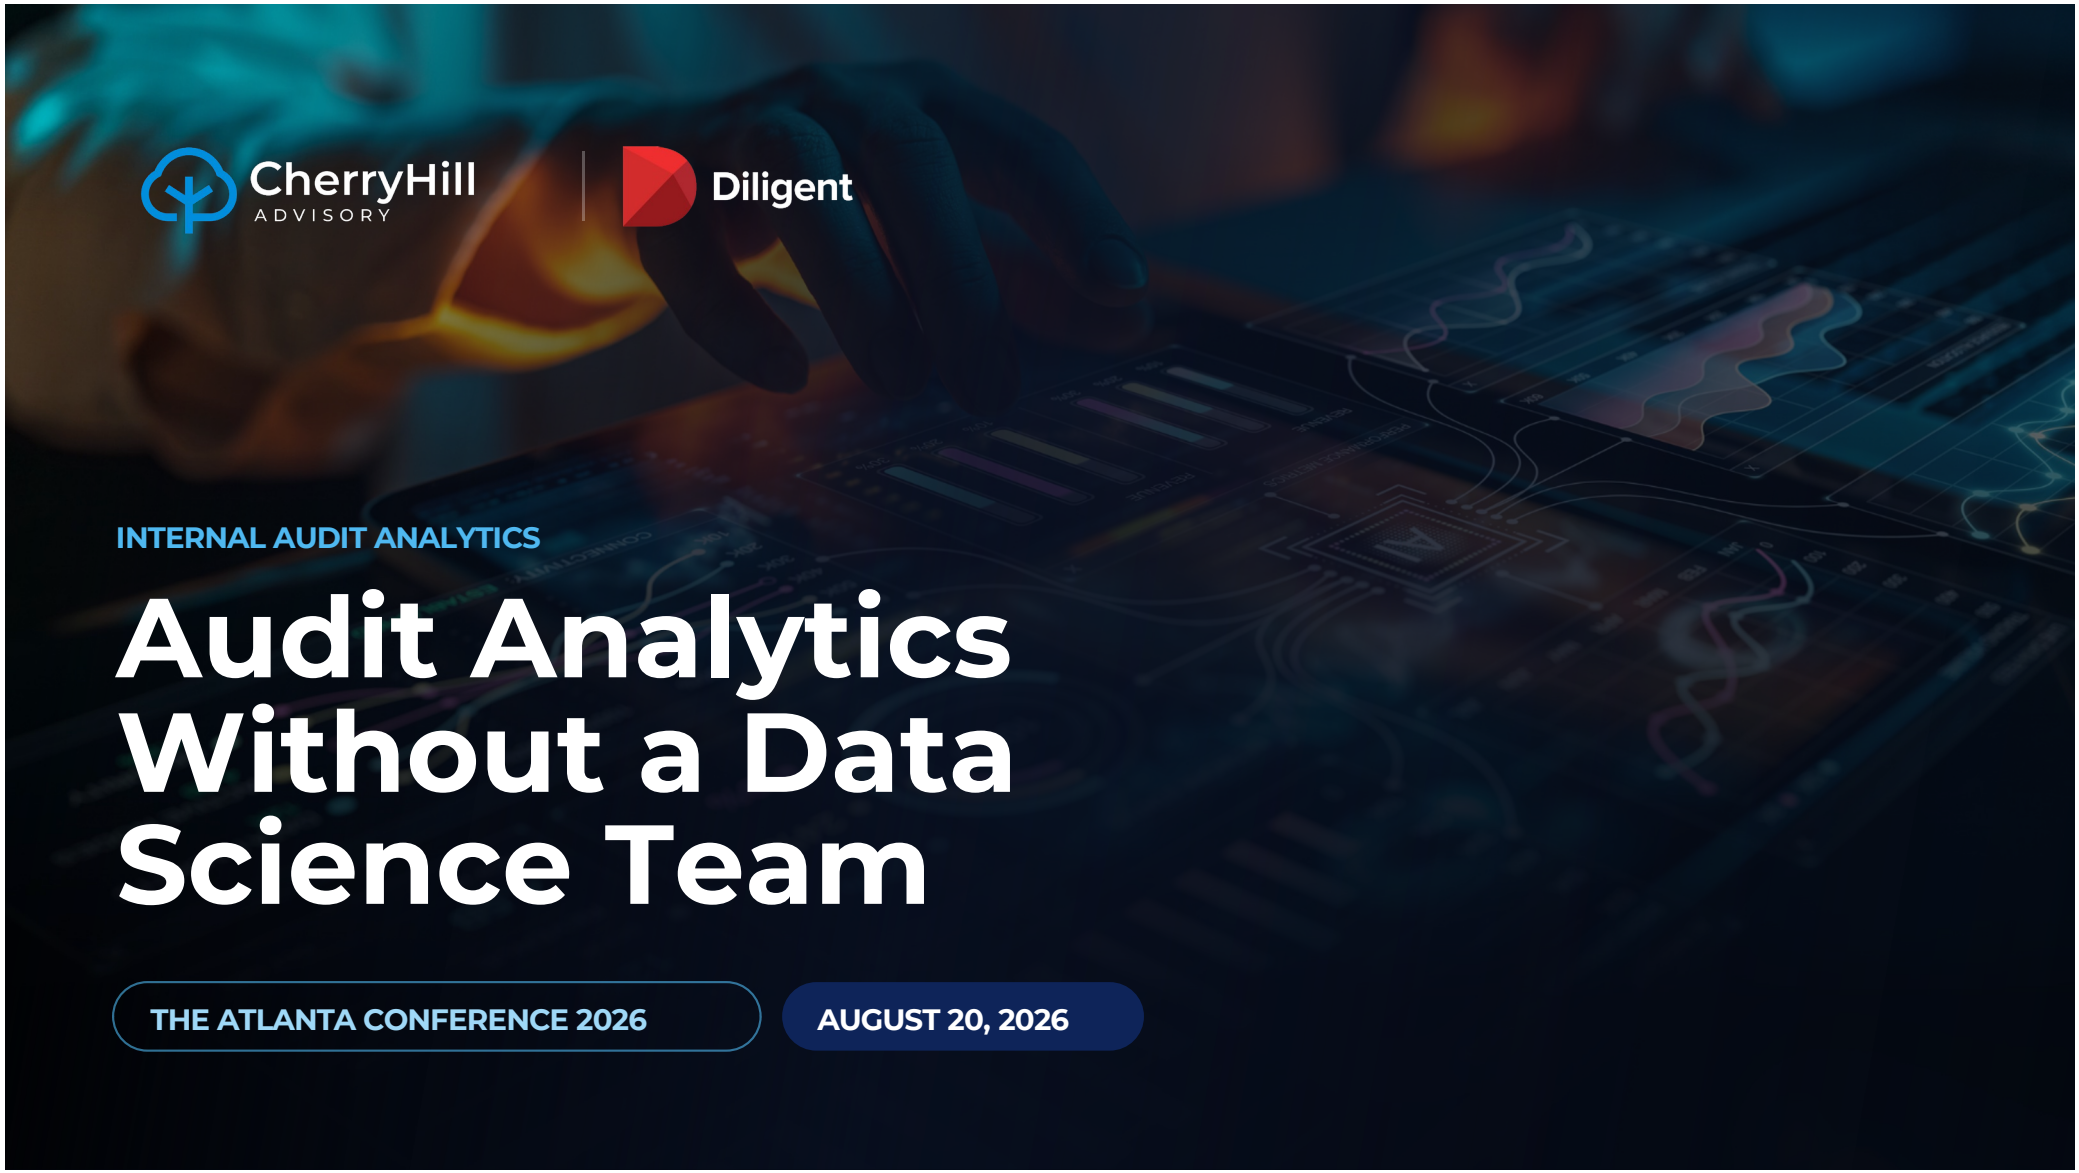
CherryHill
ADVISORY
Diligent
INTERNAL AUDIT ANALYTICS
Audit Analytics Without a Data Science Team
THE ATLANTA CONFERENCE 2026
AUGUST 20, 2026

Mike Levy

CIA, CRMA, CISA, CISSP, CDPSE, MBA

CEO and Managing Principal, Cherry Hill Advisory

Diligent Ambassador

- 20+ years in internal audit, cybersecurity, and risk management
- Former Chief Audit Executive and Deloitte consultant
- IIA International Internal Audit Standards Board member
- Past Chair, IIA North American Board; Global Board Director

CONNECT ON LINKEDIN

Learning Objectives

FOUR OBJECTIVES

1

Identify the current state and common failure points in internal audit analytics today, and recognize the conditions, data access, data quality, process stability, repeatability, that most often stall an analytics program before it starts.

2

Differentiate rule-based analytics, AI-assisted analytics, and agentic workflows across a three-tier model, and determine which tier is appropriate for a given audit objective and risk level.

3

Apply a repeatable method for translating an audit objective into a specific, testable use case, demonstrated through worked examples showing the source data, the prompt, and the output, to determine which use cases to pursue first.

4

Construct the human-in-the-loop review, documentation, and supervision steps required to rely on AI-assisted or agent-produced output as audit evidence, including what a reviewer must be able to reproduce.

Agenda

01 **Analytics Current State and Failure Points**

02 **Agentic Execution and Human Oversight**

03 **Use Cases**

01

Where Internal Audit Analytics Stands Today

Before we talk about what AI changes, here's what hasn't changed yet.

What Analytics Actually Looks Like Today

BELIEVED
A repeatable analytics program

REALITY
One person's personal method, not a documented one

BELIEVED
A shared, documented approach

REALITY
Knowledge that lives in one person's head, not a file anyone else can open

BELIEVED
Data access that's already solved

REALITY
A fresh negotiation almost every engagement

BELIEVED
Coverage across the audit plan

REALITY
A handful of tests, reused wherever they happen to fit

No function fails on all four at once. Most fail on one, and don't find out which until the person who built it leaves.

Source: ISACA, Internal Audit Data Analytics for Beginners (2023) | The IIA, Global Internal Audit Standards (2024)

The Capability Gap

1 in 3

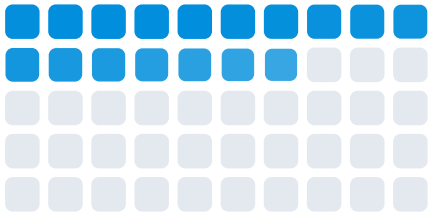

Fewer than one in three internal audit teams fully use data analytics.

92%

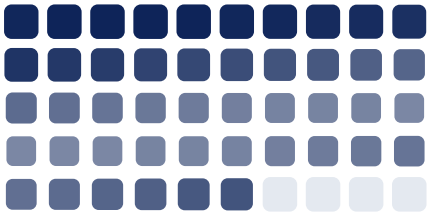

Say AI adds value. Fewer than half actually use it.

83%

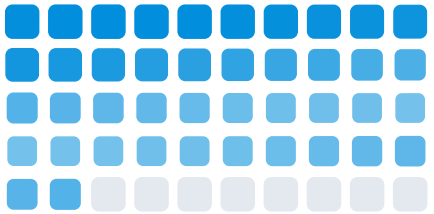

Are piloting AI, and another 12% start this year.

each square is 2% of respondents remainder

Source: Cherry Hill Advisory, Focus 2026 | Cherry Hill Advisory, The Future Internal Audit Leader (2026) | Gartner, Top Audit Department Priorities for 2026 | Jefferson Wells, Internal Audit Priorities Survey Report (2024)

REAL-WORLD EXAMPLE

How Analytics Transformed a Function, Start to Finish

2019–2020

ADNOC's Group Audit team built a Central Follow-up Analytics Solution, rolled out group-wide.

2020

A dedicated analytics team formed, and auditors were cross-trained to do the work themselves.

Day one

Governance built in from the start, not retrofitted later.

The point isn't the tooling. It's that analytics became a competency the team owned.

Source: ADNOC, "Building an Audit Intelligence Framework," Dataiku Community (2022)

Where Analytics Programs Stall

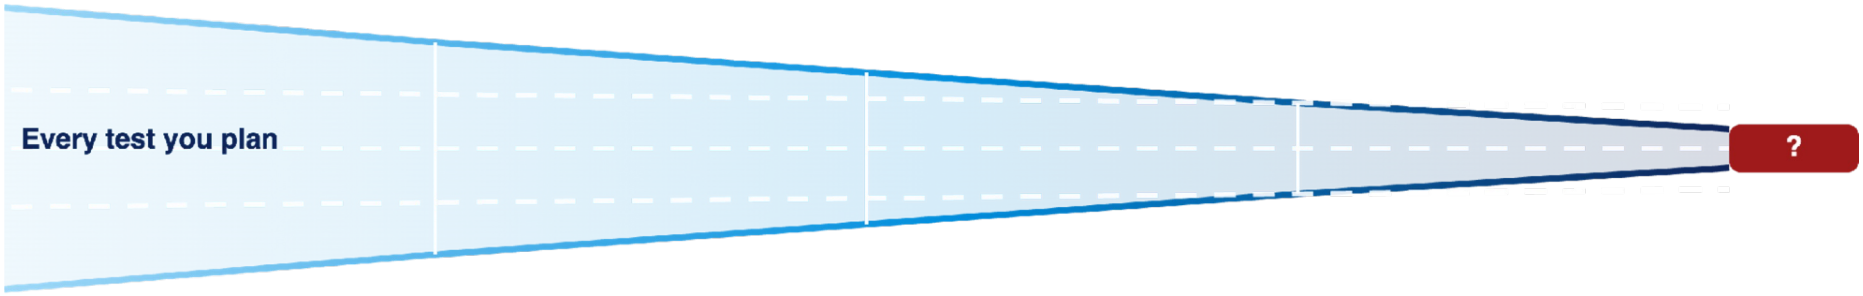

Data Access
Someone has to grant it. That person has other priorities.

Data Quality
Silent gaps. You don't find out until a result looks wrong.

Process Stability
The control changes. The test doesn't.

Repeatability
One person built it. Only one person can run it.

Still running
Whatever survives a year later.

Repeatability is the one that actually kills programs. It's invisible until the person who built the test changes roles, and by then the fix isn't "write it down," it's "rebuild from scratch," mid-engagement.

Source: ISACA, Internal Audit Data Analytics for Beginners (2023) | The IIA, Global Internal Audit Standards (2024)

Worked Scenario: The One That Actually Happens

THE TEST → WHAT HAPPENS NEXT

1 **An analyst builds a strong duplicate-payment test**

↓

Runs cleanly for 3 engagements

2 **She takes a new role in Q4**

↓

The match logic, the exclusions, the manual cleanup steps, all live only in her head

3 **The next audit needs the same test**

↓

No documentation to hand off

Result: the test gets skipped, or rebuilt from scratch under deadline, without her judgment calls baked back in.

Beyond the Four Conditions: Why Programs Actually Die

The four conditions where analytics programs stall are technical. These aren't, and they kill programs just as often:

1

No executive sponsorship.

The program exists because one manager cares. When that manager moves on, so does the mandate.

2

One owner, no backup.

The person who built it is also the only one who can explain it, defend it, or fix it when it breaks.

3

Adopted, then abandoned.

A tool gets purchased, piloted once, and never becomes part of the actual audit cycle.

4

No connection to the audit plan.

Analytics runs as a side project, separate from the objectives it's supposed to be testing.

None of these show up in a data quality check. All of them show up in an org chart.

Sequencing Your Program

01

Pick high-value, low-effort use cases first, and evaluate your own environment.

02

Prove your validation habit on lower-stakes work before applying it somewhere the stakes are higher.

03

Rescore periodically. What's low-effort today may not stay that way as systems change.

02

What Is an Agent, and How Do They Work?

What Is an Agent?

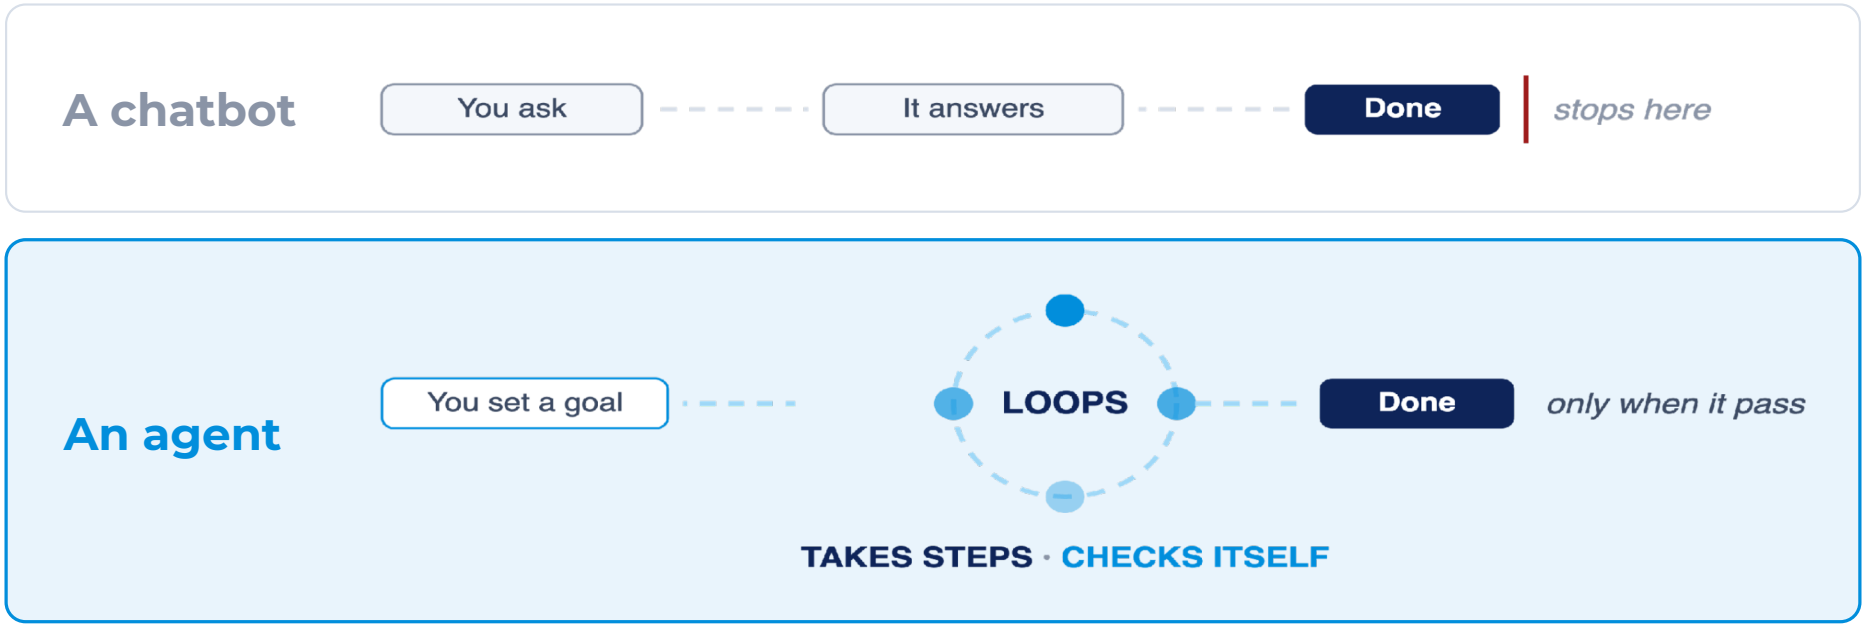

EXAMPLE "Summarize this contract" is a chatbot task. "Review this vendor's compliance and flag anything unusual" is an agent task.

REAL-WORLD EXAMPLE

Where Agentic Audit Actually Is Today

80% of tasks automated *vendor marketing claim*

☑ What agents actually do today	✖ What still requires a human
Route and follow up on requests	Set materiality
Match records against each other	Decide the scope of an engagement
Flag exceptions for review	Conclude on control effectiveness
Draft a first-pass summary	Sign off on a finding

Source: DSG.AI, "AI Agents for Internal Audit: What Actually Works in Production" (2026)

No production deployment today, from any vendor, does the right column without a human.

How Do You Use Agents in Analytics?

ONE AT A TIME

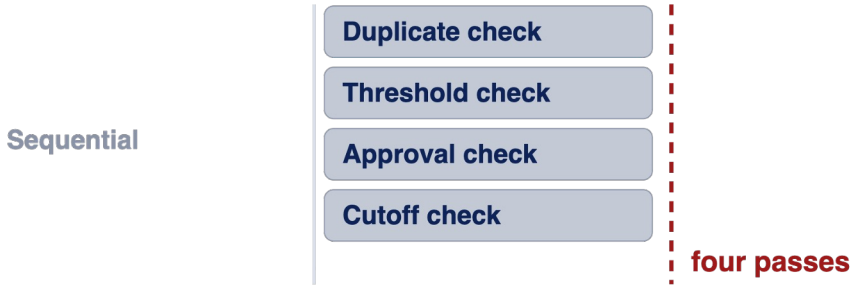

One instruction, one test.

ALL AT ONCE

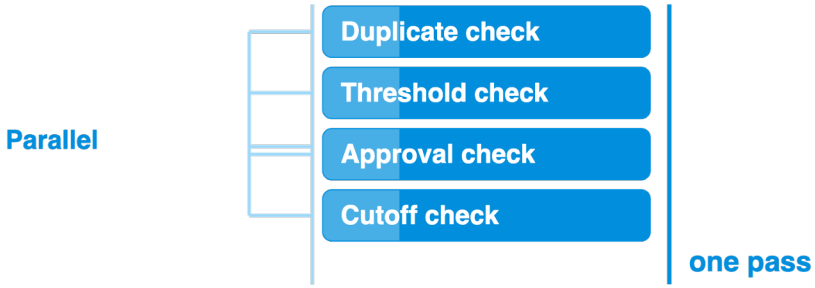

One instruction, several tests, run at once.

EXAMPLE **4×** the wall-clock difference. "Review this AP population for exceptions" produces a duplicate check, a threshold check, and an approval check, all from that single ask.

Where Do Agents Come Up, Outside Analytics?

Analytics

Running the tests themselves.

Evidence Requests

Chasing down PBC (Prepared By Client) items automatically.

Risk-Control Mapping

Matching a new risk to the controls that address it.

Project Setup

Scaffolding a new audit from scope to timeline.

Agents in Practice: Evidence Requests

Turn evidence collection from chasing into coordinated execution

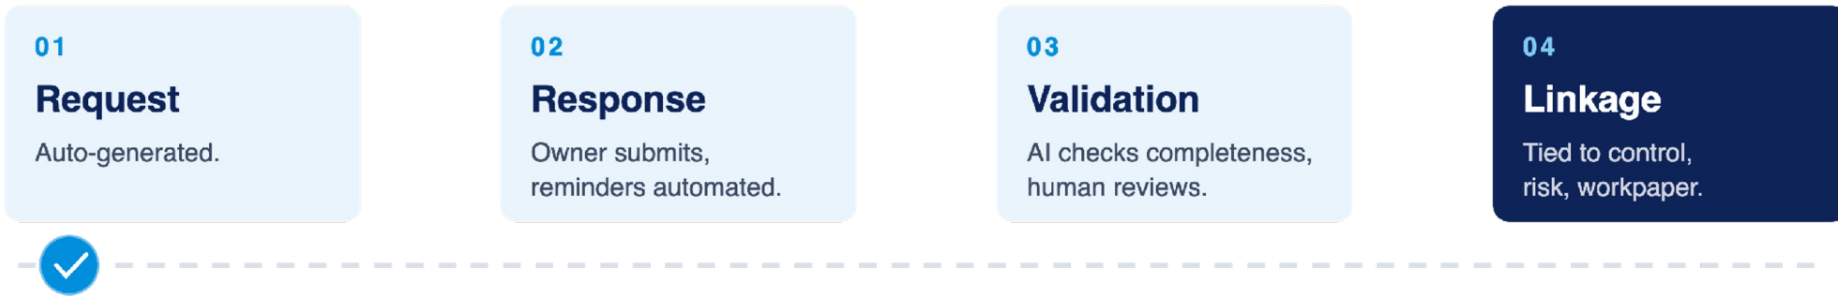

EXAMPLE An auditor needs a bank confirmation from Treasury. The agent drafts and sends the request, tracks who hasn't responded, checks completeness, and links it to the control, without anyone chasing an email thread.

Source: 2026 IIA Singapore Presentation (Diligent)

Agents in a Platform vs. Agents in an Large Language Model (LLM) Product

Built in, running quietly. You didn't have to ask for it.

You invoke it directly. You ask, and it tells you it's running several agents at once on your request.

Quick Tips: How to Actually Invoke an Agent

1 **Give it a goal, not a question.**

"What's a duplicate payment?" gets an answer. "Review this AP population for exceptions" gets an agent doing work.

2 **Name the specific checks.**

Don't say "find problems," say what to look for, same vendor and amount within 30 days, missing approvals, whatever the actual test is.

3 **Ask for one pass, not three.**

If you need duplicates, thresholds, and approvals checked, ask for all three in one instruction, that's the whole point of an agent over a single-answer chatbot.

4 **Ask it to show its work.**

Request the reasoning behind each flag, not just a final list, that's what you'll actually be reviewing.

5 **Review before you trust it.**

An agent finishing a task isn't the same as the task being right.

The AI Maturity Curve

share of the work done by people done by AI

The next era of internal audit is defined by AI that does more than assist.

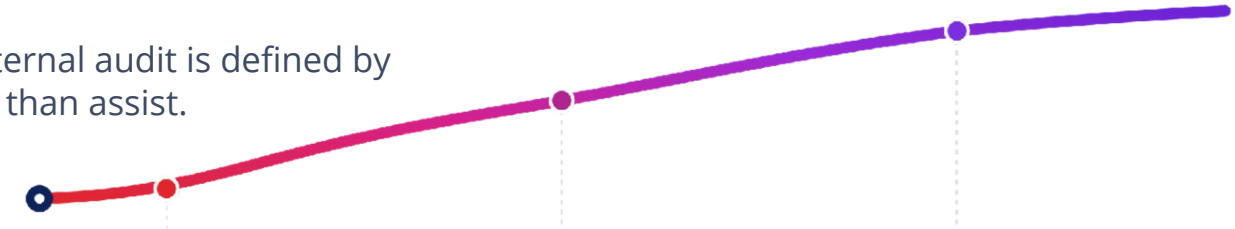

LEVEL 1

AI Assistants

Embed AI into existing workflows. Less about transformation, more about productivity, summarizing information, identifying risks faster.

LEVEL 2 · HUMAN IN THE LOOP

Native AI

AI isn't just assisting, it's actually doing work, handling tasks, driving workflows.

LEVEL 3 · HUMAN ON THE LOOP

Autonomous Agents

AI orchestrates workflows end-to-end. Humans provide oversight rather than execution.

Source: 2026 IIA Singapore Presentation (Diligent)

Build vs. Buy: The Key Decisions

	Build It Yourself	Buy a Platform
Pricing	Tool subscription only, your time is the real cost	Platform subscription, vendor absorbs the build cost
Time to stand up	Days to weeks, depends on your prompting skill	Faster, the workflow is pre-built
Subject matter expertise	Your team owns and maintains the knowledge	The vendor carries the expertise for you
Consistency of output	Requires deliberate configuration to avoid hallucinations	Built-in logic is designed to run the same way every time
Learning curve	Requires learning to prompt effectively	Largely button-click, less prompt skill needed
Data handling and audit trail	You're responsible for logging every prompt and output yourself	Typically built on enterprise data-handling terms with logging as a baseline feature
Ease of deployment across a team	Depends on teaching others your exact approach	The same workflow is available to your whole team immediately

FROM EARLY SIGNAL TO AUTONOMOUS ACTION

Either Way, One Question Stays With You

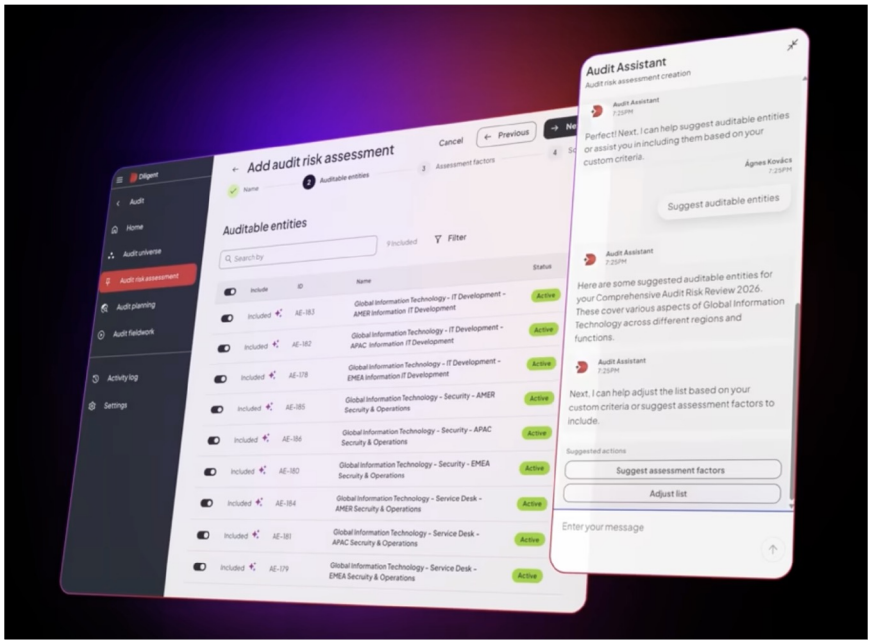

Who validates the output, and can you prove it.

~70%

reduction in request management time

~$10k

estimated savings per audit

SmartPrep 360 detects emerging risk signals at the board level. AuditAI converts those signals into dynamic audit scope updates, intelligent evidence requests, and structured remediation tracking, creating continuous, risk-aligned assurance.

03

Use Cases

The method, then six worked examples, each showing the problem, the source data, the prompt, and the output.

EXAMPLE 1 OF 6 · DUPLICATE PAYMENT DETECTION

PROBLEM / SOURCE DATA / PROMPT / OUTPUT

The Problem

Invoice `#100804` for $17,436.26. Seventeen days later, the same vendor submits `#100805`, same amount. Maybe a legitimate resubmission, maybe the same payment processed twice by two people who never checked each other's work.

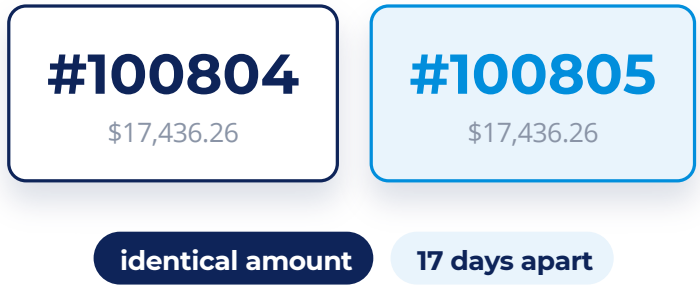

EXAMPLE 1 OF 6 · DUPLICATE PAYMENT DETECTION

The Source Data

AP TRANSACTION ROWS

Invoice	Vendor	Amount	Date
INV-100000	Northstar Logistics	$1,271.73	21 May 2025
INV-100003	Meridian Facilities Group	$1,486.41	22 Apr 2025
INV-100804	Coastal Freight Solutions	$17,436.26	13 Oct 2025

THE MATCHED PAIR, CLOSE UP

100804 · Coastal Freight Solutions · $17,436.26 · 13 Oct

100805 · Coastal Freight Solutions · $17,436.26 · 30 Oct

VENDOR FIELD, ISOLATED

Coastal Freight Solutions

Coastal Freight Solutions

Exact match

DATE FIELD, 17-DAY GAP

13 Oct — **30 Oct**

Inside the 30-day test window.

Population: 840 AP transactions, one fiscal year. *No payment date or check number in the file.*

The Prompt

ROLE	Act as a senior internal auditor with expertise in accounts payable and procurement controls.
CONTEXT	I am conducting an audit of the accounts payable process. The scope covers 840 AP transactions from one fiscal year.
TASK	Identify potential duplicate payments. Look for: same vendor and amount within 30 days, same invoice number, similar but not identical invoice numbers.
FORMAT	A table showing vendor, both dates, amount, days apart. Rank by dollar value.
CONSTRAINTS	Professional audit language. Flag assumptions. Do not conclude fraud, only flag for investigation.

EXAMPLE 1 OF 6 · DUPLICATE PAYMENT DETECTION

Actual Output

Accounts Payable — Duplicate Payment Exception Analysis

Item	Detail
Population tested	840 AP transactions (source file: example1_ap_transactions.csv)
Period covered	2025-01-01 through 2025-12-30 (one fiscal year, per client representation)
Total value of population	$17,781,088.69
Distinct vendors	10
Data fields available	invoice_number, vendor, amount, invoice_date
Test 1 — Same vendor + identical amount within 30 days	18 pairs / 36 transactions
Test 2 — Same vendor + identical amount beyond 30 days	2 pairs / 4 transactions
Test 3 — Identical invoice number (exact)	0 exceptions
Test 4 — Similar (non-identical) invoice numbers	Not evaluable — see 'Test 3-4 Invoice Nos' tab
Test 5 — Identical amount, different vendor	1 pair (see 'Other Observations' tab)
Test 6 — Near-identical amounts, same vendor, <=30 days	57 pairs (low precision — see 'Other Observations')

Assumptions and Limitations

1. The data set contains no payment date, payment/check number, payment status, purchase order number, invoice line detail, remittance reference, or credit-memo indicator. Testing is therefore performed on INVOICE records, not on confirmed disbursements. An exception below may represent a duplicate invoice that was blocked, offset by a credit, or never paid.

2. 'Vendor' is treated as an exact string match. No vendor master file was provided, so related-party or duplicate vendor records (same supplier under two vendor IDs, DBA names, or spelling variants) cannot be tested. The 10 vendor names in the file are all distinct with no apparent aliasing.

See the full output

Scan for all 20 pairs, the full prompt, and the assumptions

EXAMPLE 1 OF 6 · DUPLICATE PAYMENT DETECTION

Output and Analysis

Rank	Vendor	Date 1	Date 2	Amount	Days Apart
1	Coastal Freight Solutions	13 Oct	30 Oct	$17,436.26	17 days
2	Meridian Facilities Group	07 Apr	03 May	$15,788.36	26 days
3	Pinecrest Office Supply	30 Jan	17 Feb	$14,601.08	18 days
—	15 further pairs	—	—	—	—

18 pairs
flagged for investigation

$161,196
total matched value

1 of 3
requested tests was a dead end

ANALYSIS The invoice numbers are sequential, so "similar invoice numbers" was a dead test, the model said so instead of inventing a list.

Source: Cherry Hill Advisory methodology, run against a test population to verify the approach.

EXAMPLE 2 OF 6 · REVENUE RECOGNITION TREND ANALYSIS

The Problem

A single month's revenue number means nothing on its own. It only means something against a baseline, and building that baseline across five regions and thirty-six months is work nobody has time for.

So most teams don't, and the movement that mattered stays invisible.

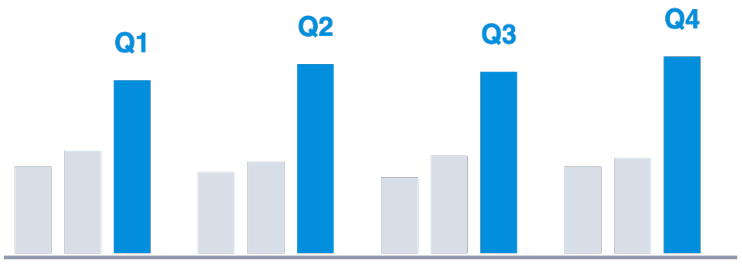

EXAMPLE 2 OF 6 · REVENUE RECOGNITION TREND ANALYSIS

The Source Data

SAMPLE ROWS

Region	Month	Revenue
Southeast	Jan 2023	$376,104.07
Southeast	Feb 2023	$357,336.98
Southeast	Mar 2023	$379,009.03

SOUTHEAST'S FULL 36-MONTH SEQUENCE

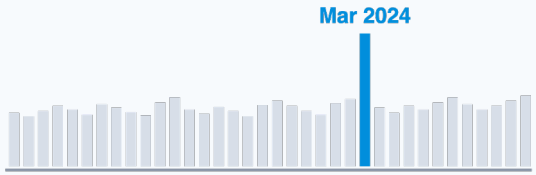

WHAT THE EYE MISSES

177 values:*2 decimals*

3 values: *4 decimals*

Population: 180 region-months across 5 regions and 3 years · 165 testable

The Prompt

ROLE	Act as an audit data analyst specializing in revenue recognition risk.
CONTEXT	Reviewing revenue trends across 5 regions over 3 years, 180 data points.
TASK	Identify trends month-over-month and quarter-over-quarter. Flag any month with movement over 15% from the trailing 3-month average. Run Benford's Law on the same figures.
FORMAT	Table of flagged months by region and percentage deviation.
CONSTRAINTS	No fraud conclusions. Distinguish known drivers from unexplained flags.

Actual Output

	A	B	C	D	E	F	G	H	I
1	**EXCEPTIONS - MONTHS DEVIATING MORE THAN 15% FROM TRAILING 3-MONTH AVERAGE**								
2	*6 exceptions from 165 scoped observations (3.6%). Sorted by absolute deviation, descending.*								
3									
4	**Region**	**Month**	**Revenue**	**Trailing 3-Mo Avg**	**Deviation**	**MoM %**	**Direction**	**Classification**	**Basis / Analyst Note**
5	Southeast	2024-03	$508,518	$368,909	+37.8%	+35.0%	Increase	Unexplained	Largest positive deviation in the population, 4.7x the next-largest March deviation anywhere in the population (+8.0%, Midwest 2025). March is seasonally strong (index 1.09), but Southeast Mar-2025 deviated only +3.3%; Mar-2023 has no trailing window. Value carries 4 decimal places. Also contaminates the trailing window for Apr-Jun 2024.
6	West	2025-04	$290,272	$383,910	-24.4%	-26.0%	Decrease	Unexplained	Largest single-month contraction in the population. West April was flat against its trailing window in both prior years (+1.0% in 2023, +0.0% in 2024); April carries a seasonal index of 1.03, i.e. a normal-to-strong month. Value carries 4 decimal places, unlike every clean observation in the series. Requires GL and cut-off corroboration.
7	Midwest	2025-10	$392,167	$322,960	+21.4%	+20.3%	Increase	Unexplained	October is seasonally soft (index 0.97); this is a +21% expansion against that pattern. Midwest Oct-2023 and Oct-2024 deviated -1.3% and +1.6%. Value carries 4 decimal places.
8	Southeast	2025-02	$410,579	$346,542	+18.5%	+10.5%	Increase	Partially explained	Q1 is the seasonal peak across all regions (Feb index 1.06, Mar 1.09) and Q1 follows the Q4-to-Q1 step-up seen in all three years. Magnitude is at the upper edge of that pattern but directionally consistent.
9	Southeast	2024-06	$348,207	$417,093	-16.5%	-4.3%	Decrease	Derived / known driver	Artifact of the Mar-2024 outlier inflating the Apr-Jun trailing window. Substituting the Feb/Apr-2024 mean for March moves this from -16.5% to -6.8%, below threshold. Not an independent exception.
	Southwest	2025-12	$289,609	$350,817	+15.5%	+13.8%	Increase	Partially explained	December rebound is a recurring pattern in every region and year (Dec index 1.02 against a Sep-Nov

README | Monthly Detail | Flag Summary | Watch Tier 10-15% | QoQ Trend | Seasonality | Benford | Data Integrity Note

See the full output

Scan to download The full workpaper: all 165 tested months, the seasonality index, the Benford test, and the prompt.

EXAMPLE 2 OF 6 · REVENUE RECOGNITION TREND ANALYSIS

Output and Analysis

Region	Month	Move vs. trailing avg.	Assessment
Southeast	Mar 2024	+37.8%	Known driver, promotion
West	Apr 2025	-24.4%	Known driver, contract renegotiation
Midwest	Oct 2025	+21.4%	Unexplained, referred
3 further months	—	—	Dismissed as noise

165
observations tested

6
exceeded the threshold

3
genuinely unexplained

ANALYSIS One of the six wasn't real. March's outlier sat inside June's trailing window, and removing it puts June at −6.8%, below threshold. No tool flags that. Reading the six as a set does.

EXAMPLE 3 OF 6 · USER ACCESS TERMINATION TESTING

The Problem

An employee is terminated Friday. The access-removal ticket sits three days in a backed-up IT queue. For those three days, someone with no reason to be in the building still has a working login.

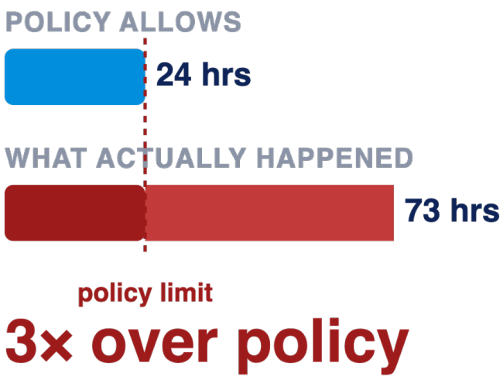

- **What policy allows** 24 hours
- **What actually happened** 73 hours, three times over

The removal ticket sat in a backed-up queue across a weekend. Nobody did anything wrong, and the login worked the whole time.

The Source Data

SOURCE ROWS

ID	Terminated	Disabled	Gap
E-10450	20 Aug 2025	Same day	19.1 hrs
E-10451	27 Aug 2025	Same day	10.0 hrs

THE FLAGGED EXCEPTION, CLOSE UP

E-10513 · terminated 14 Feb 2025

disabled 17 Feb 2025, 1:00 AM

73-hour gap

A CLEAN EXAMPLE FOR CONTRAST

E-10452 · terminated 28 Oct 2025

Disabled 4.3 hours later

Population: 64 terminations, full year

The Prompt

ROLE	Act as an IT auditor specializing in access management controls.
CONTEXT	Testing whether system access was disabled within 24 hours of termination, full population.
TASK	Match termination dates against disable timestamps. Flag any gap over 24 hours.
FORMAT	Table: employee ID, termination date, disable date, gap, flagged yes/no.
CONSTRAINTS	Full population, not a sample. Flag missing data.

EXAMPLE 3 OF 6 · USER ACCESS TERMINATION TESTING

Actual Output

	A	B	C	D	E	F
1	**User Access Termination Testing — Control Test Summary**					
2						
3		**CONTROL TESTED**				
4		System access is disabled within 24 hours of employee termination.				
5						
6		**SCOPE AND RESULTS**				
7		Terminations in population			**64**	
8		Terminations tested			**64**	
9		Coverage			**100%**	
10		Exceptions — gap over 24 hours			**1**	
11		Exception rate			**1.6%**	
12						
13		**GAP DISTRIBUTION (HOURS)**				
14		Shortest gap			**0.4**	
15		Median gap			**9.6**	
16		Average gap			**11.2**	
17		Longest gap within policy			**21.9**	
18		Longest gap overall			**73.0**	
19		Cleared the limit by under 3 hours (21.0–24.0)			**5**	
20						
21		**DATA QUALITY CHECKS**				
22		Records with a missing field			**0**	
23		Duplicate employee IDs			**0**	

Summary | Test Results | Methodology

See the full output

Scan to download all 64 tested terminations and the full methodology.

EXAMPLE 3 OF 6 · USER ACCESS TERMINATION TESTING

PROBLEM / SOURCE DATA / PROMPT / **OUTPUT**

Output and Analysis

Employee	Gap vs. 24-hour policy	Flagged
E-10513	73 hours	Yes
63 further terminations	0.4 – 21.9 hrs	No

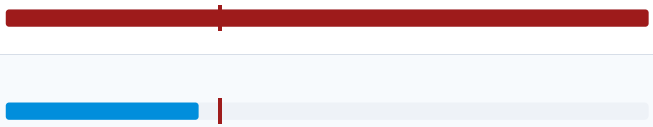

1
exception out of 64 tested

73 hours (3 days)
3× the 24-hour policy limit

0
missing data records

ANALYSIS Testing the full population, not a sample, is what makes this finding possible, a 40-item sample had a 37% chance of missing the one employee this actually happened to.

EXAMPLE 4 OF 6 · SEGREGATION OF DUTIES AND PRIVILEGED ACCESS REVIEW

The Problem

Creating a vendor and approving a payment are each perfectly reasonable permissions on their own. The risk only exists when the same person holds both, and a 120-row access list will never surface that by reading down it.

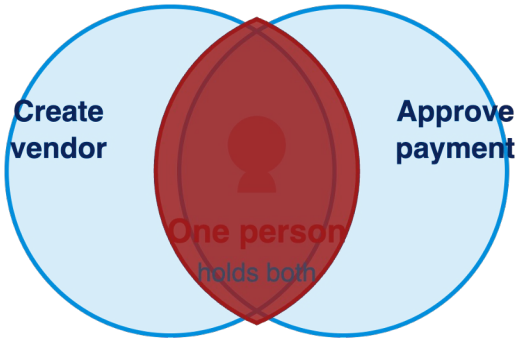

Fine apart. A conflict only in the overlap.

EXAMPLE 4 OF 6 · SEGREGATION OF DUTIES AND PRIVILEGED ACCESS REVIEW

The Source Data

SAMPLE ROWS

user_id	assigned_roles
M-1005	Purchase Requisition; Journal Entry; Reconciliation
M-1007	PO Approval; Payment Approval; Purchase Requisition; Vendor Maintenance
M-1086	PO Approval; Payment Approval; General Ledger Access; Payroll Approval; Reconciliation
M-1090	Read-Only Reporting

...116 more rows

11

distinct role names across the population, no spelling variants, no fuzzy matching needed

- No transaction data. No dollar amounts, no dates, no HR feed.
- No system of record required, this is an export any application owner can produce in ten minutes.

Population: 120 users

EXAMPLE 4 OF 6 · SEGREGATION OF DUTIES AND PRIVILEGED ACCESS REVIEW

The Prompt

ROLE	Act as an IT auditor specializing in segregation of duties.
CONTEXT	Full user population, testing for incompatible duty combinations.
TASK	Identify SOD conflicts: vendor maintenance + payment approval; journal entry + reconciliation; purchase requisition + PO approval; payroll processing + payroll approval. Also identify any user holding 3 or more high-risk roles simultaneously (general ledger access, payroll approval, PO approval, payment approval).
FORMAT	Conflict matrix by user, type, risk.
CONSTRAINTS	Cross-reference role combinations carefully before flagging

Actual Output

Conflict Matrix - by User, Conflict Type and Risk

User ID	Rule	Conflict Type	Conflicting Roles Held	Risk Rating	All Roles Assigned	User Composite Risk	Findings for User	Risk Rationale	Additional Observation
M-1005	C2	Journal Entry + Reconciliation	Journal Entry; Reconciliation	High	Purchase Requisition; Journal Entry; Reconciliation	High	1	Records journal entries and performs the reconciliation designed to detect them. Self-review threat; permits	
M-1006	C3	Purchase Requisition + PO Approval	Purchase Requisition; PO Approval	Moderate	PO Approval; Purchase Requisition	Moderate	1	Initiates and approves the same purchase order. Unauthorized or personal-benefit procurement;	
M-1007	C1	Vendor Maintenance + Payment Approval	Vendor Maintenance; Payment Approval	High	PO Approval; Payment Approval; Purchase Requisition; Vendor	High	2	Can create or modify a vendor master record and approve payment against it. Fictitious-vendor and	End-to-end procure-to-pay exposure: holds 4 of 4 P2P roles.
M-1007	C3	Purchase Requisition + PO Approval	Purchase Requisition; PO Approval	Moderate	PO Approval; Payment Approval; Purchase Requisition; Vendor	High	2	Initiates and approves the same purchase order. Unauthorized or personal-benefit procurement;	End-to-end procure-to-pay exposure: holds 4 of 4 P2P roles.
M-1014	C2	Journal Entry + Reconciliation	Journal Entry; Reconciliation; Bank Reconciliation	High	General Ledger Access; Payroll Approval; Reconciliation; Journal Entry;	High	1	Records journal entries and performs the reconciliation designed to detect them. Self-review threat; permits	
M-1023	C1	Vendor Maintenance + Payment Approval	Vendor Maintenance; Payment Approval	High	Payment Approval; Vendor Maintenance; Payroll Processing	High	1	Can create or modify a vendor master record and approve payment against it. Fictitious-vendor and	
M-1046	C3	Purchase Requisition + PO Approval	Purchase Requisition; PO Approval	Moderate	Purchase Requisition; PO Approval; Payroll Processing	Moderate	1	Initiates and approves the same purchase order. Unauthorized or personal-benefit procurement;	

Summary | Conflict Matrix | 1 User Role Detail

See the full output

Scan to download all 16 findings and the full population data.

EXAMPLE 4 OF 6 · SEGREGATION OF DUTIES AND PRIVILEGED ACCESS REVIEW

Output and Analysis

Conflict type	Users	Count	Risk
Requisition + PO Approval	M-1006, M-1007, M-1046, M-1048, M-1067, M-1117	6	Moderat
Vendor Maintenance + Payment Approval	M-1007, M-1023, M-1112, M-1115	4	High
Journal Entry + Reconciliation	M-1005, M-1014, M-1074, M-1096	4	High
Payroll Processing + Payroll Approval	M-1118	1	High
High-Risk Role Concentration	M-1086	1	Critical

120
users tested

15
users with a conflict · 12.5%

16
separate findings

1
critical exception

M-1086 holds all four high-risk roles. No pair rule catches it, only counting does.

M-1007 appears twice. Two findings, one account, the whole procure-to-pay cycle.

EXAMPLE 5 OF 6 · SOC 2 REPORT REVIEW

The Problem

One SOC 2 report is 40 to 80 pages. A third-party risk program reviewing twenty vendors a year isn't hard because any single review is difficult, it's hard because doing all twenty with the same rigor competes with everything else on the calendar.

1,200

pages, one reviewer

at 60 pages a report

EXAMPLE 5 OF 6 · SOC 2 REPORT REVIEW

The Source Data

This example uses a constructed, illustrative SOC 2 excerpt (Northwind Cloud Services, Type II, Security and Availability, 2025 period).

SECTIONS SHOWN

System description

Cloud storage and processing
2 of 5 trust categories examined
12-month period, Jan–Dec 2025

Control test results

3 controls shown
4.2 Logical Access Reviews
5.1 Environmental Safeguards
5.4 Incident Response

Subservice organization disclosure

1 carve-out
3 control areas excluded
Physical security, environmental, facility access

The Prompt

ROLE	Act as a third-party risk auditor.
CONTEXT	Annual vendor risk assessment, reviewing a SOC 2 Type II report.
TASK	Summarize trust service criteria covered, list exceptions noted, identify subservice organization carve-outs, flag control gaps.
FORMAT	Structured summary with risk rating and rationale.
CONSTRAINTS	Flag carve-outs for follow-up regardless of stated rationale.

EXAMPLE 5 OF 6 · SOC 2 REPORT REVIEW

Actual Output

OVERALL RISK RATING **MODERATE** provisional	Rating is provisional pending receipt of the subservice organization's SOC 2 report and written clarification of the scope inconsistency described in Section 4. Both items are prerequisites to a final rating.

1. Trust Services Criteria Covered

Category	In scope	Notes
Security (Common Criteria)	Yes	Baseline; required in all SOC 2 engagements
Availability	Yes	Examined
Confidentiality	**No**	Not examined
Processing Integrity	**No**	Not examined
Privacy	**No**	Not examined

Coverage is the minimum-plus-one configuration: two of five categories. For a vendor that stores and processes entity data, the absence of Confidentiality and Processing Integrity is a substantive scope limitation, not a formality. Those two categories are the ones most closely aligned to what the vendor actually does with our data.

See the full output

Scan for the full memo, source excerpt, and prompt.

EXAMPLE 5 OF 6 · SOC 2 REPORT REVIEW

Output and Analysis

Item	Detail
Access review timeliness	2 of 4 quarters late, 41 days (Q1), 12 days (Q3)
Scope coverage	Security and Availability only; 3 categories not examined
Subservice carve-out	Regional Data Centers LLC, physical, environmental, facility access

2 of 4
quarterly access reviews completed late

3 of 5
trust categories not examined

Moderate
provisional risk rating

ANALYSIS

The review surfaced something a skim wouldn't: the report carves out environmental controls to the subservice organization, then tests an environmental control and reports no exceptions. Those can't both be true. That contradiction, plus the carve-out itself, is what goes on the follow-up list before renewal.

EXAMPLE 6 OF 6 · THIRD-PARTY CONTRACT COMPLIANCE

The Problem

A 1.5% late penalty. A 3% annual escalation tied to CPI. Nobody re-reads the clause when an invoice lands, they check the number looks about right. Two years and two escalations later, no one has confirmed the invoice still matches the contract.

Invoiced amount

Contract clause

Year 1 → Year 2 · nobody re-checks

EXAMPLE 6 OF 6 · THIRD-PARTY CONTRACT COMPLIANCE

The Source Data

> A constructed, illustrative professional services agreement excerpt, managed IT support, $175,000 base monthly fee. Nine clauses across three sections.

SECTIONS SHOWN

Fees and Payment Sec. 3.1–3.4	Service Levels Sec. 5.1–5.3	Termination Sec. 8.1–8.2
`4 clauses`	`3 clauses`	`2 clauses`
$175,000 monthly base	99.5% monthly availability	90 days for convenience
$200–$350 hourly, out of scope	4-hour critical response	Either party, no cause needed
1.5% per month late penalty	5% monthly service credit	30-day cure on material breach
3% "tied to CPI" escalation		

The Prompt

ROLE	Act as an internal auditor reviewing contract compliance.
CONTEXT	Professional services agreement review.
TASK	Summarize payment terms, obligations, SLAs. Flag unusual terms creating audit risk.
FORMAT	Contract review memo, 5-7 recommended audit procedures.
CONSTRAINTS	Flag ambiguous terms rather than assume standard practice.

EXAMPLE 6 OF 6 THIRD-PARTY CONTRACT COMPLIANCE

Actual Output

5. Flagged Terms and Audit Risk

The following provisions are ambiguous, internally inconsistent, or asymmetric in a way that creates measurable audit or financial exposure. Each is stated as written; Internal Audit has not assumed a market-standard reading.

Ref	Provision	Issue identified	Risk
F-1	§3.4 Escalation	The clause is internally inconsistent. It fixes the increase at 3% and simultaneously ties it to the Consumer Price Index, which are different mechanics that produce different results. The clause does not name the CPI series (CPI-U, CPI-W, national or regional), the measurement period, the publication used, or whether 3% operates as a cap, a floor, or the sole rate. A one-point difference between the fixed rate and an index-driven rate changes the five-year spend by approximately $225,000.	High
F-2	§3.2 Rate band	Out-of-scope work is billed within a $200–$350 range "depending on staff seniority," a 75% spread. The Agreement defines no seniority tiers, no rate card mapping title to rate, no requirement for written pre-approval of out-of-scope work, and no cap on aggregate out-of-scope billings. Rate selection rests with the Vendor. This is the single largest uncontrolled cost exposure in the Agreement.	High

See the full output

Scan for the full 7-page memo, 9 flagged terms and 7 audit procedures.

EXAMPLE 6 OF 6 · THIRD-PARTY CONTRACT COMPLIANCE

Output and Analysis

Item	Detail
Payment terms	$175,000/month, plus $200–$350/hour uncapped
Escalation Sec. 3.4	"3%" and "CPI-tied", two different mechanics, one clause
SLA remedy Sec. 5.3	Flat 5% credit, same for a 0.1-point miss or a month-long outage
Termination Sec. 8.1	90 days, either party, no transition obligation

9
terms flagged as audit risk

7
audit procedures recommended

$225K
5-year exposure from one ambiguous clause

ANALYSIS Section 3.4 fixes the increase at 3% and ties it to CPI, those are different mechanics producing different numbers. Nobody catches that reading an invoice. The memo turns nine ambiguous terms into seven testable procedures.

Agents Running in Parallel

Each of these six worked independently today. In practice, they don't have to run one at a time.

Key Takeaways

01 Start with a testable question, not a tool.

02 Six worked use cases, each with a real prompt and a documented output.

03 None of this requires a data science team, it requires a method and a validation habit.

Continue the Conversation

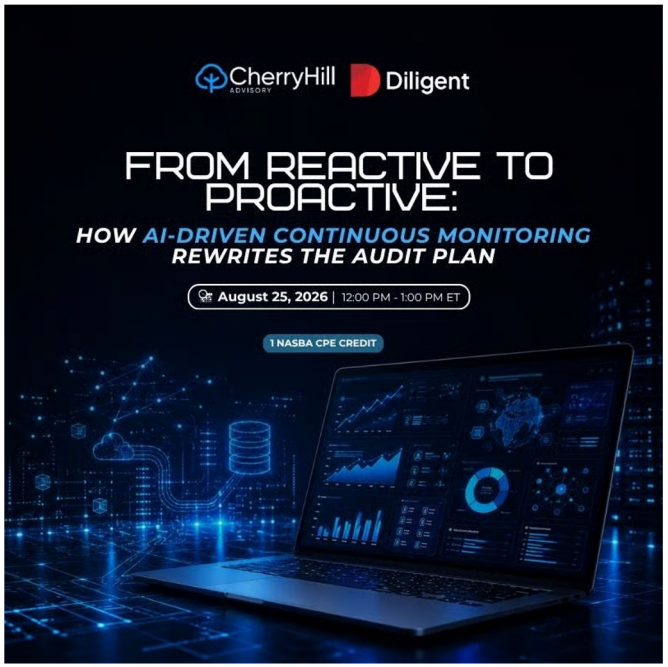

Scan to register

1 NASBA CPE credit per session.

Resources for Today

AI Internal Audit Encyclopedia

The Future Internal Audit Leader

Focus 2026: Setting the Internal Audit Agenda

The 2026 Pulse of Internal Audit, CHA Risk Register

IIA 2026 North American Pulse of Internal Audit Survey

Q&A

Mike Levy, CIA, CRMA, CISA, CISSP, CDPSE, MBA

CEO and Managing Principal, Cherry Hill Advisory

Connect with me on LinkedIn