Internal Audit Collective

Build, Buy or Both?
Putting AI to Work in Internal Audit

Internal Audit Collective

November 17, 2026
www.InternalAuditCollective.com

Internal Audit Collective

Engagement Requirements

Engage Actively:
Participate actively by responding to all polling questions.

Technical Preparedness:
Ensure your device is set up with a stable internet connection, functioning audio, and video capabilities if required.

Stay Engaged:
To receive the credited hour, you must remain on the call for the entire duration of the webinar.

Provide Feedback:
Complete the post-webinar surveys to share your insights and improve future sessions.

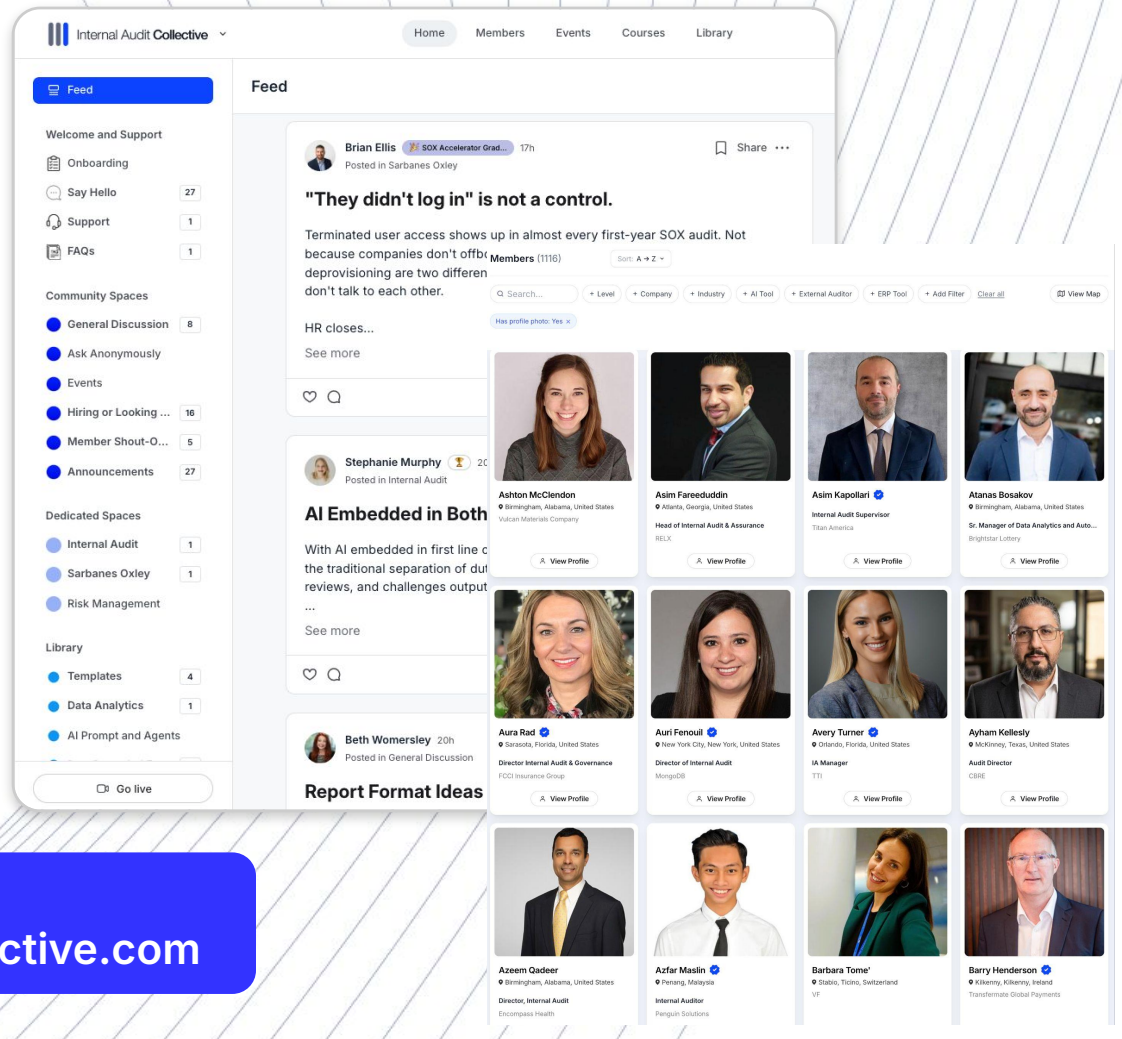
Join the Internal Audit Community Collective Today

Scan to join the Collective
Use code **WEB1117** for 20% off
www.InternalAuditCollective.com

Join an online, "managed" community for SOX and Internal Audit professionals

The Collective is a membership site designed for SOX and Internal Auditors to network, share, learn, and grow.

Get access to an online resource to ask questions, get feedback, and exchange ideas with your peers.



Join today at: www.InternalAuditCollective.com

Take the full framework back to your team



Inside: both scorecards, about 40 audit activities mapped to tools, and five case studies of teams that built, bought or did both.

By David Coulombe and Tom O'Reilly. Co-written with the Internal Audit Collective, a community and training organization for SOX and internal audit professionals.

Scan for the eBook

petual.ai/resources/build-or-buy-ai-strategy

Or ask for a copy at booth 15.

petual

Two lenses on one decision: audit and procurement



David Coulombe

Audit Executive, Petual

VP of Internal Audit, Lyft

The internal audit view: what the work has to meet

Built and scaled Lyft's first internal audit function, after leading global audit teams at Intel, Juniper Networks and Salesforce, and starting at EY in Paris.

Co-author of the Internal Audit Collective eBook on build or buy. CIA, and a board member of the IIA San Francisco and San Jose chapters.

Controls Director, Intel

The procurement view: what it takes to buy and run it

Owned the policies and processes behind Intel's end-to-end procurement cycle: supplier selection, contracting and performance monitoring.

Covered the bulk of Intel's procurement activity worldwide.

petual

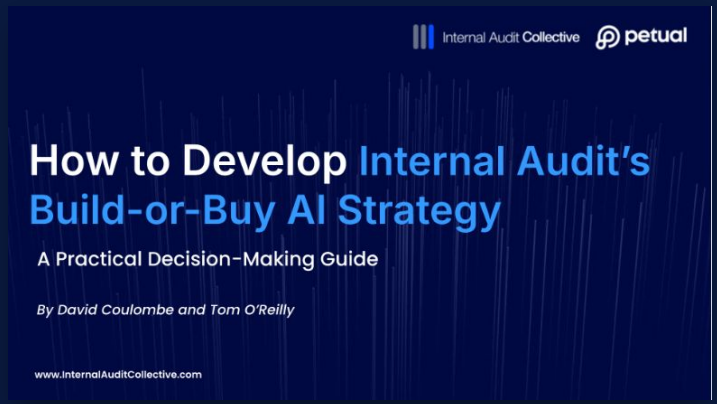
Internal Audit Collective

LIVE POLL 1

Where is your internal audit team today with AI?

1. Not yet using it
2. Individuals are experimenting
3. Piloting specific use cases as a team
4. Using it regularly in established workflows

FREE EBOOK FROM IAC



Scan to download

AI is here to stay. Our job now is paving the road

The question for internal audit is no longer whether AI becomes part of the work. It is how: which work it takes on, how we rely on it, and who builds it.

petual

AI will change how our teams operate at every layer, so it belongs at the core of the IA strategy

1

More coverage from the same team, as long as the quality bar holds.

2

New skills, a changing reviewer role, new staffing choices.

3

The business is making the same AI calls, and we need to understand those firsthand.

petual

LLMs can do more every month, so building looks easy. Is it?

petual

Build or buy is a continuum, and each use case can sit at a different point on it

Buy — **Hybrid** — **Build**

Adopt a vendor-developed AI solution.

The vendor owns most product development and maintenance.

Configure purchased AI platforms to your methodology, and build bounded workflows or agents on top.

Responsibility is shared between the vendor's platform and your internal team.

Develop and maintain your own AI applications, agents, integrations and orchestration.

Your organization owns most of the engineering, testing, governance and maintenance.

petual

Decide in three steps, starting from the work, not the tool

1

Classify the work by task type

Seven types, from research and drafting through to executing a test and concluding.

2

Match the type of AI tool that fits

Each task type has a tool best suited to carry it. Use the lightest one that does the job reliably.

3

Decide what you can realistically build

And maintain, as processes, data and models change. Then decide build or buy, one use case at a time.

petual

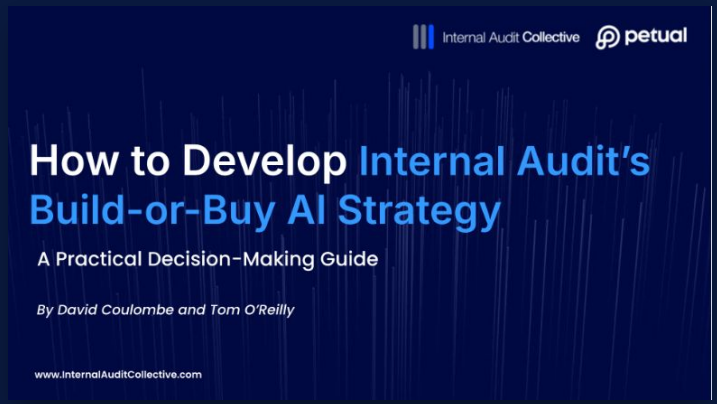
Internal Audit Collective

LIVE POLL 2

If you could meaningfully improve one part of your audit work with AI, which would you choose?

1. Research, risk assessment, and planning
2. Evidence requests and follow-ups
3. Evidence analysis and control testing
4. Findings, reporting, and communications

FREE EBOOK FROM IAC



Scan to download

Classify the work by the nature of the task

AI assists the auditor

Learn
Research and summarize. *Lease accounting under GAAP.*

Brainstorm
Generate risks, objectives, scope, questions.

Build
RCMs, work programs, templates, scripts.

Write
Emails, findings, reports. *First drafts in one team style.*

AI performs part of the work

Coordinate
Requests, reminders, scheduling, status, escalations.

Analyze
Extract and compare documents and data against criteria.

Execute and conclude
Gather evidence, test, document, conclude. *Invoice-to-PO reconciliation.*

petual

Three kinds of AI tool, all running on models you buy

Less autonomy and reach → More autonomy and reach

LLM assistant

A frontier model in an enterprise chat interface.

The person: drives every step and owns the output.

Looks like: enterprise chat and office-suite assistants.

Simple agent

One defined workflow with a few tool connections.

The person: triggers it and reviews the output.

Looks like: custom agents in an approved builder.

Complex agentic system

Orchestrated agents working across many systems, end to end.

The person: sets scope, reviews and signs off.

Looks like: purpose-built audit platforms or engineering programs.

The frontier model layer sits under all three. Everyone buys it; nobody here builds it.

petual

A simple agent in practice: regulatory change sensing

Trigger

Scheduled task

Weekly, Monday 6am. Runs even with your laptop closed.

Sources

Connectors

Web search for SEC, PCAOB and regulator feeds; a connector for paid news. Read-only.

Agent

Skill and files

Written steps, a scoring rubric and a digest template. Audit universe and risk register as project files.

Output

Microsoft 365

Digest as a Word doc, sent by Outlook. Write access set by your M365 admin.

Human

Reviewer

The audit manager edits and approves. Claude writes to neither the register nor the plan.

What the Monday task says

"Every Monday at 6am, check the SEC, PCAOB and state regulator feeds and the industry alerts for anything published since last week. Compare each item to the audit universe and risk register. Score relevance high, medium or low, draft the one-page digest with proposed audit plan impacts, and send it to the audit team. Do not change the risk register or the audit plan."

Governance, mapped to the tool

Source allow-list: approved connectors only, read-only.

Run log: every run kept as its own session.

Quarterly accuracy check: sample past digests, score hits and misses, tune the steps.

petual

Models are powerful, but the harness is the difference.

The tools on the last slide could share the same models. A harness is where the specific needs for the application live

The harness

Instructions
A model improvises. Instructions make it follow the same steps every time.

Tools and data
A model only knows what it's shown. This decides what it can see and touch.

Memory
A model forgets between runs. This keeps corrections and clarifications.

Foundation model

Raw capability.
Can't check its own work.

Orchestration
A model answers one prompt. This turns many steps into one piece of work.

Checks
A model can't tell when it's wrong. Checks catch it and send the work back.

Guardrails and logs
A model leaves no trail. This limits what it may touch and records what it did.

In a different metaphor, the model is the engine, but there's more to the car . Steering, brakes, the chasis, and instruments are what make it safe to rely on and use, and that's where the build effort goes.

petual

What an audit harness gives you

Work that will be relied on has to be traceable, explainable and reperformable, whoever or whatever performed it.

The standard you already hold

Traceable

Every conclusion points to the exact evidence it rests on.

Tools and data · Guardrails and logs

Explainable

The reasoning is written down, step by step, in terms a reviewer can follow.

Instructions · Memory

Reperformable

Another auditor can repeat the test from the documentation alone.

Instructions · Guardrails and logs

What it takes to rely on it at scale

Consistent

Same steps on every sample, whoever runs it. Same evidence, same conclusion.

Instructions · Orchestration

Contained

It sees only what it should, and won't otherwise invent to bridge a gap

Tools and data · Guardrails and logs

Measurable

Scored against known answers, so you see when a change makes it worse.

Checks · Guardrails and logs

The model doesn't make work traceable, explainable or reperformable. The harness does.

petual

What it takes to build one yourself

What the system has to do, on every sample

Extract evidence → **Map to samples** → **Perform the test** → **Generate documentation**

What it takes to build and run one in-house

Engineers

Several, dedicated, who understand both the models and the audit work.

After launch: The team stays. The system is never finished.

Data access

Connections into your ERP, GRC tool and evidence stores.

After launch: Kept working as those systems change.

Infrastructure

Secure hosting, model APIs, document processing and storage for evidence.

After launch: Patched, monitored and paid for every month.

Evals

Large test sets with known answers, and people to review the results by hand.

After launch: Re-scored every time a model or prompt changes.

Governance

Security review, access controls, logging and data-handling agreements.

After launch: Revisited with every change, and every audit of the tool.

Maintenance

Fixes, new evidence formats, and changes to the controls you test.

After launch: Models change every few months. Each one means re-testing.

All very doable, but not necessarily simple

petual

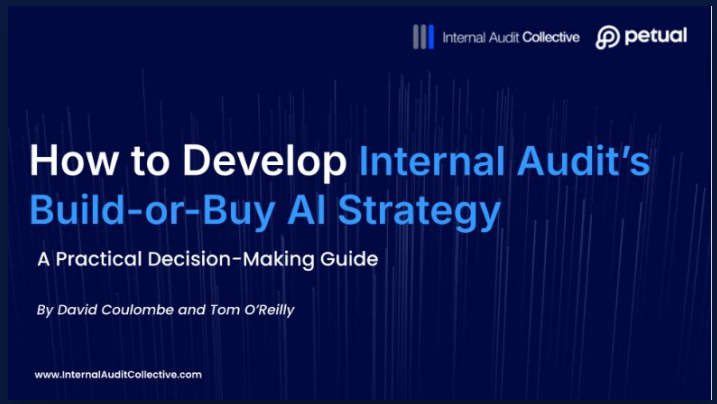
Internal Audit Collective

LIVE POLL 3

What would you most need to see before relying on AI-generated control testing?

1. Strong results against cases our team has independently evaluated
2. Consistent results when the same test is repeated
3. Conclusions clearly linked to evidence, with work someone else can re-perform
4. Clear review responsibilities and agreement on acceptable use

FREE EBOOK FROM IAC



Scan to download

Built or bought, test it the same way

1

Take a few controls your team tested by hand last year and run them through the tool on the same samples. Keep the disagreements.

2

Run them again with nothing changed. If an answer moves, you want to know before your auditor does.

3

Hand the workpapers to your reviewer. Would they sign them, and could someone else re-perform them?

Next: which of your work actually needs a system like this, and which doesn’t?

Match each task type to the lightest tool that reliably does it

Task nature	What it covers	LLM assistant	Simple agent	Complex agentic
Learn	Research and summarize (e.g., regulations, standards, processes, technologies, methodology, prior audit history)	Primary	Supporting	Supporting
Brainstorm	Generate options: risks, audit objectives, scope, recommendations, questions to ask	Primary		
Build	Create reusable artifacts: RCMs, work programs, templates, questionnaires, analytics scripts	Primary	Supporting	
Write	Draft and refine communications: emails, memos, findings, reports, committee materials	Primary	Supporting	
Coordinate	Move work between people and systems: requests, reminders, scheduling, status tracking, escalations	Supporting	Primary	Supporting
Analyze	Extract and compare information against criteria: documents, policies, contracts, datasets, issue logs	Supporting	Supporting	Primary
Execute & Conclude	Gather evidence, perform the test, document it, and reach a conclusion that will be relied upon		Supporting	Primary

Building is easiest at the top. The bottom two rows are primarily fieldwork, and that's where building gets hard.

Primary best suited to the task · **Supporting** useful for part of the task · blank: limited fit · Source: the eBook, p. 22

petual

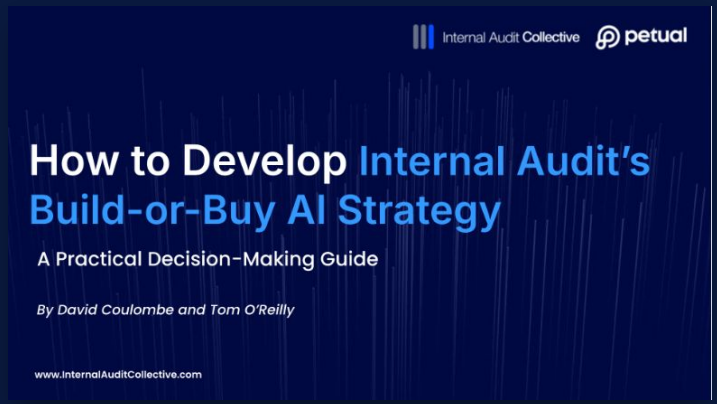
Internal Audit Collective

LIVE POLL 4

For an AI use case your team wants to pursue, what would most influence whether you build, buy, or do both?

1. Total cost to implement and maintain
2. Quality and reliability of the output
3. Governance, security, and accountability
4. Internal skills and capacity to support it
5. Fit with our methodology, systems, and scope

FREE EBOOK FROM IAC



Scan to download



Decide build or buy one use case at a time, against your team’s objective

1. **Objective:** set once for the function. It decides how you weigh everything that follows.
2. **Gate:** removes build when there is no case for owning the capability.
3. **Five dimensions:** applied to each use case, weighted to your objective.
4. **Score the path:** vendor selection if buy survives, build metrics if build survives. Both scorecards are in the eBook.

petual

Five questions to ask of every build or buy option

Dimension	The question	What sits under it
Cost	What will it cost over its whole life, including review time and the cost of getting it wrong?	Usage costs that scale, continued tuning, monitoring and drift detection
Quality	Will its output meet audit standards, and say so when the evidence isn't there?	Fabrication risk, consistency and reproducibility, performance, drift over time
Accountability	Who owns each conclusion, and can we trace it back to its evidence?	Model transparency, AI data governance, risk-based human review, audit trail, AI inventory
Skill	Do we have the people to run it, and what happens when the person who built it leaves?	Configuration skill, the reviewer's changing job, key-person risk, learning after go-live
Scope	What will it realistically cover, and how far can it stretch?	Breadth across control types, automation ceiling, portability, coverage

Ask any tool, bought or built, to trace one conclusion to its evidence, explain it, and let someone re-perform it.

petual

Take the full framework back to your team



Inside: both scorecards, about 40 audit activities mapped to tools, and five case studies of teams that built, bought or did both.

By David Coulombe and Tom O'Reilly. Co-written with the Internal Audit Collective, a community and training organization for SOX and internal audit professionals.

Scan for the eBook

petual.ai/resources/build-or-buy-ai-strategy

petual

“Our objective has not changed, but the best way to achieve it has.”

On building an audit management system, then buying a better one

Set the objective once

What you want AI to enable, and the outcomes that matter to your function.

Decide per use case

By the nature of the task, the tool it needs, and what you can really build and maintain.

Keep asking

Is our answer still the right answer? As the tools change, it may not be.

petual

Thank you

Questions are welcome.



David Coulombe
Audit Executive, Petual
david@petual.ai



Tom O'Reilly
Founder, Internal Audit Collective
tom@internalauditcollective.com

AI Office Hours

IA Internal Audit SOX SOX

Every session 12:00 – 1:00 PM ET · Live on Zoom

TUE **NOV 17**	IA	**Root Cause Analysis**	with **Brianna Bazzy** 
TUE **DEC 1**	IA	**Issue / Exception Drafting**	with **Ian Burnett** 
WED **DEC 2**	SOX	**IPE Review**	with **Patrick Noll** 
TUE **DEC 8**	SOX	**SOX TOC Logical Access**	with **Dawn Vogel** 
TUE **DEC 15**	IA	**Audit Committee Reporting**	with **Trent Russell** 
WED **DEC 16**	IA	**Issue Remediation**	with **Haseeb Hassan Khan** 

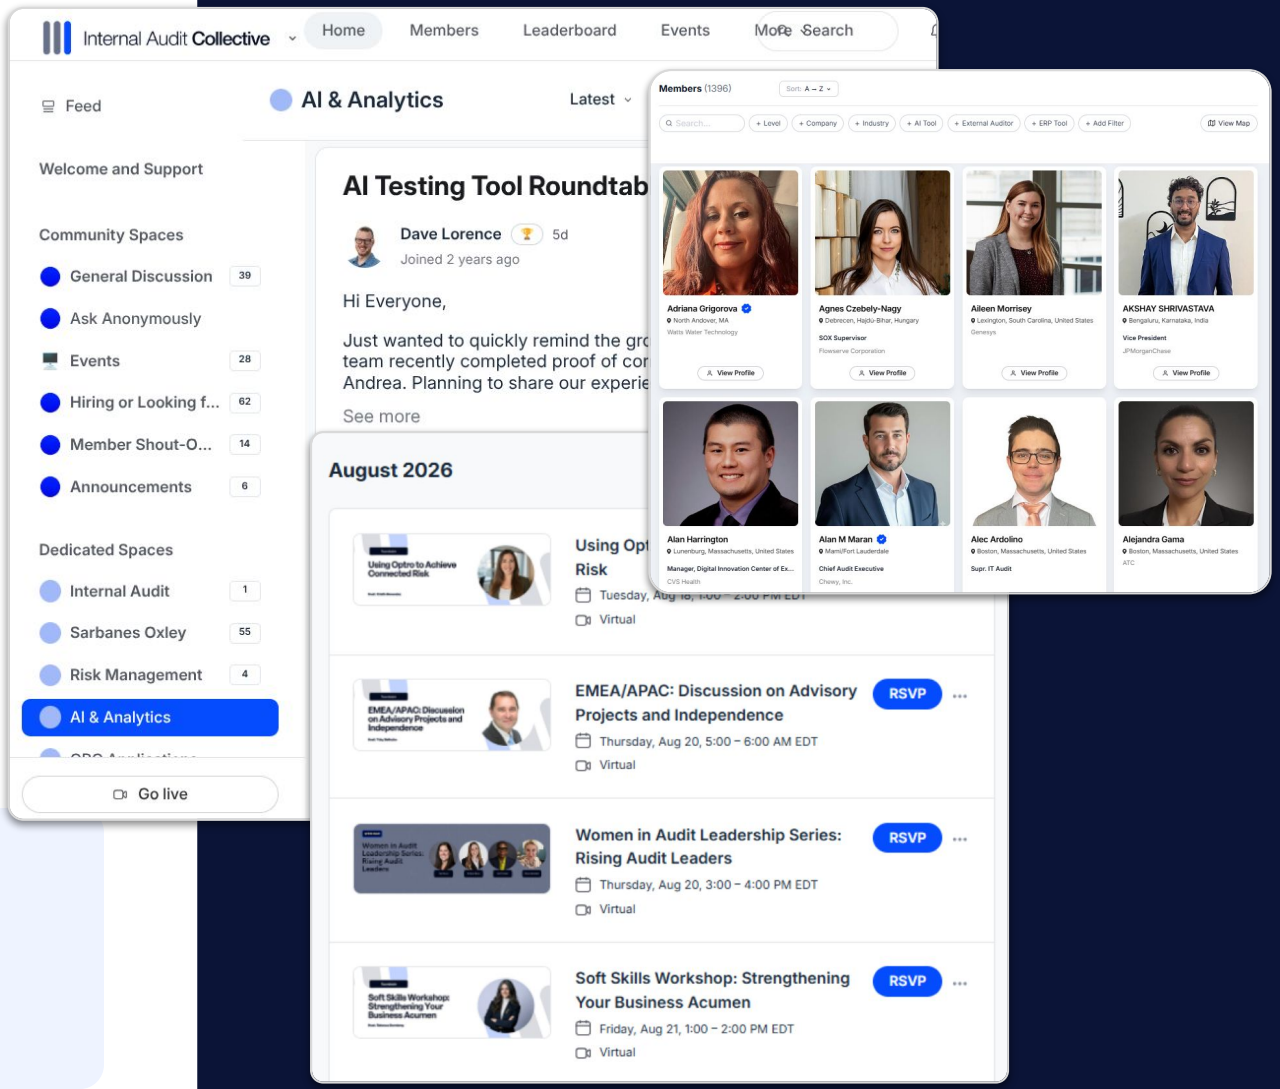
RSVP in the Events space · community.internalauditcollective.com

INTERNAL AUDIT COLLECTIVE

Join the Internal Audit Collective Community

No pitches. No spam. Just community.

- **A managed community**
  for SOX and Internal Audit professionals
- **Network, share, learn, and grow**
  with practitioners in a membership site built for you
- **Ask questions and get feedback**
  and exchange ideas with your peers



Scan to join the Collective

Use code **WEB1117** for 20% off

www.InternalAuditCollective.com

One control, end to end

The complex agentic system from the last slide, shown in the one I know best.

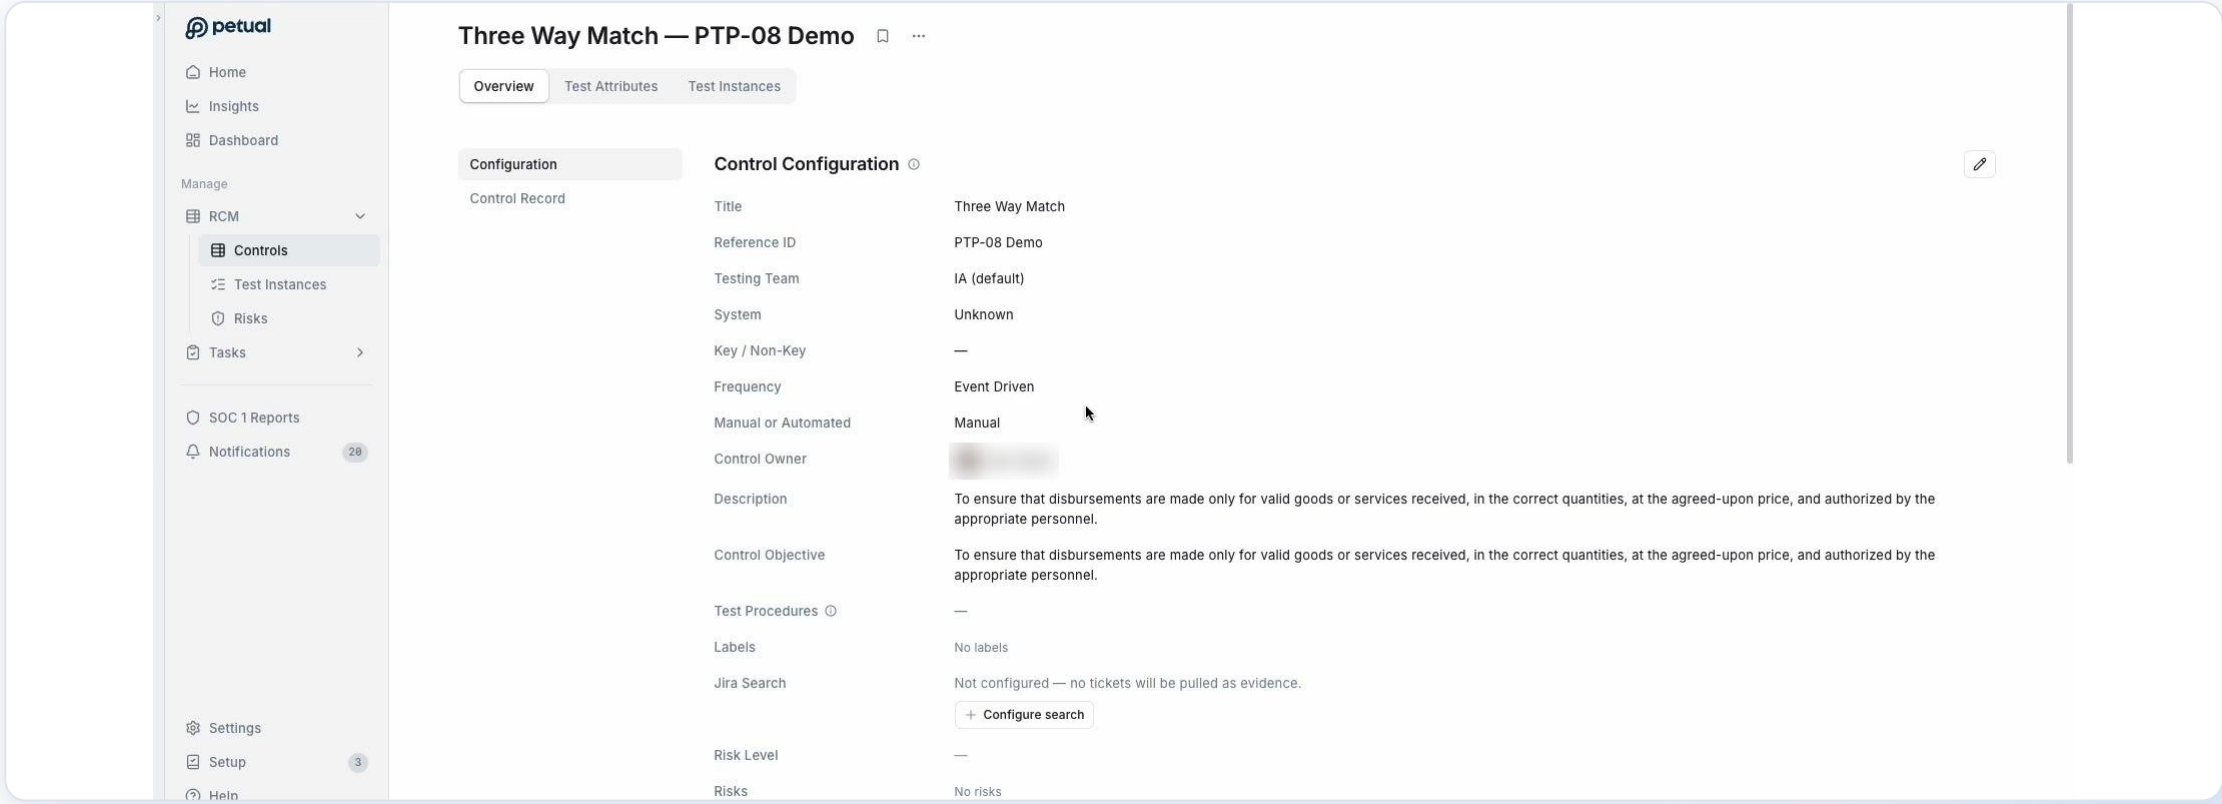
Whoever builds it, this is the work it has to do.



Petual, in a sandbox with test data: a three-way match on purchase orders.

1 Control | 2 Attributes | 3 Evidence | 4 Samples | 5 Conclusion | 6 Workpaper

Start from one control in your RCM: a three-way match

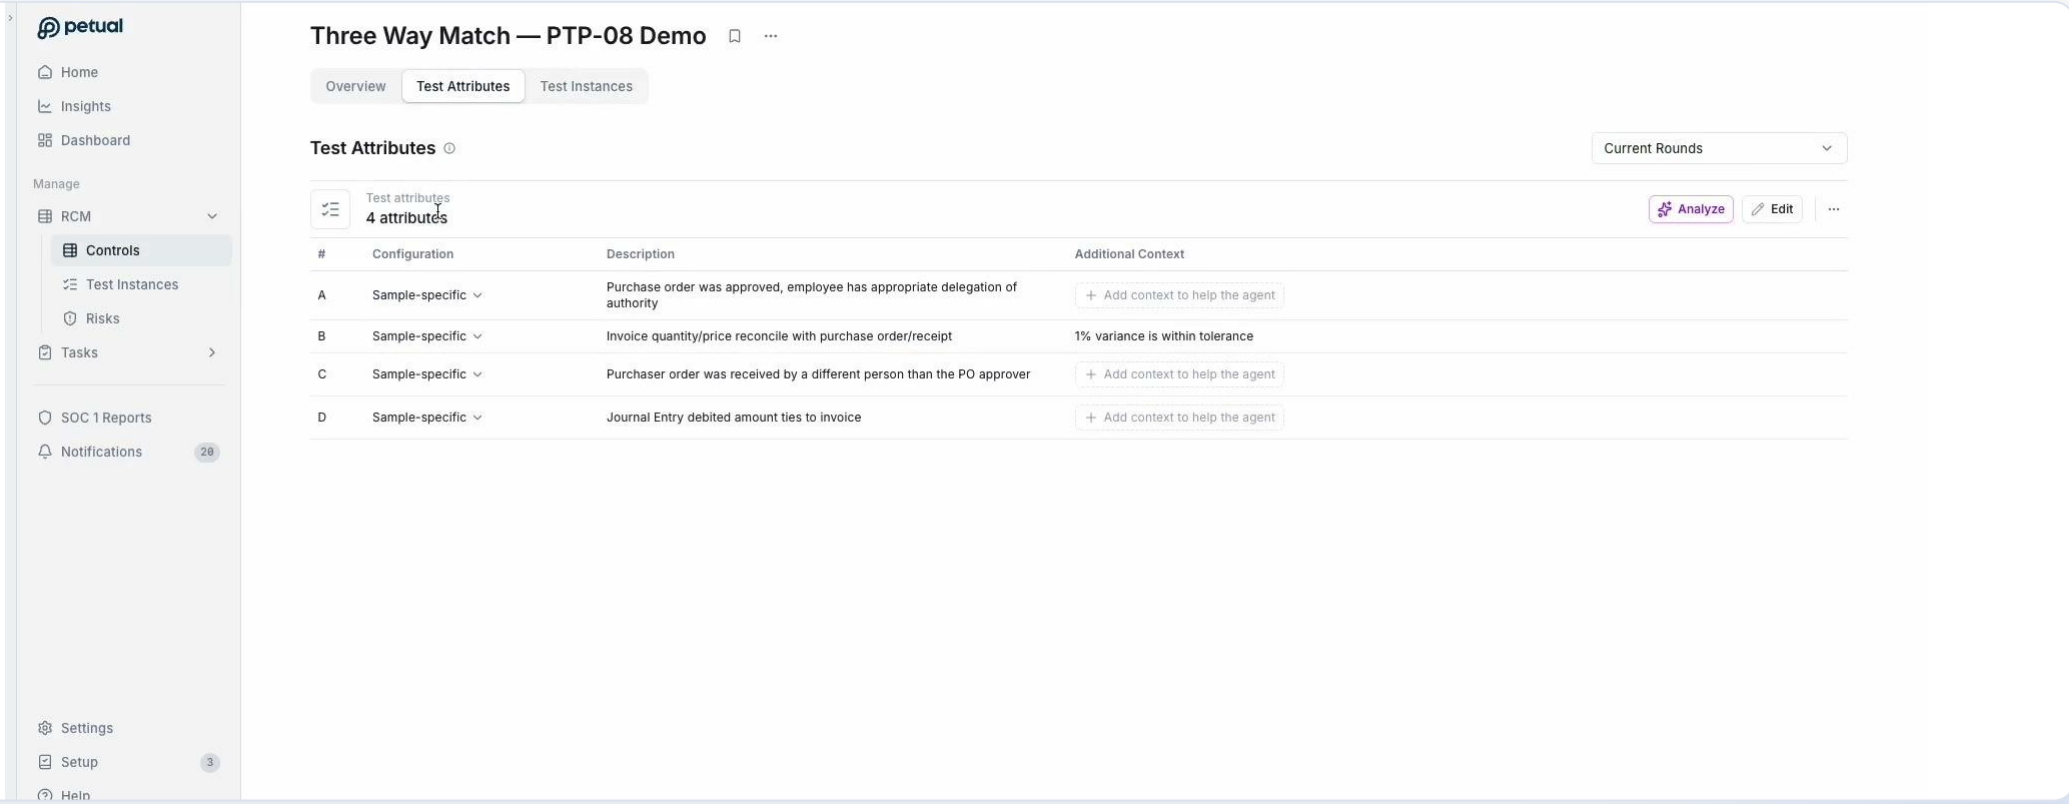


petual

1 Control | **2 Attributes** | 3 Evidence | 4 Samples | 5 Conclusion | 6 Workpaper

Harness at work · Instructions
Each attribute becomes explicit test steps.

Break it into test attributes, each checked for testability

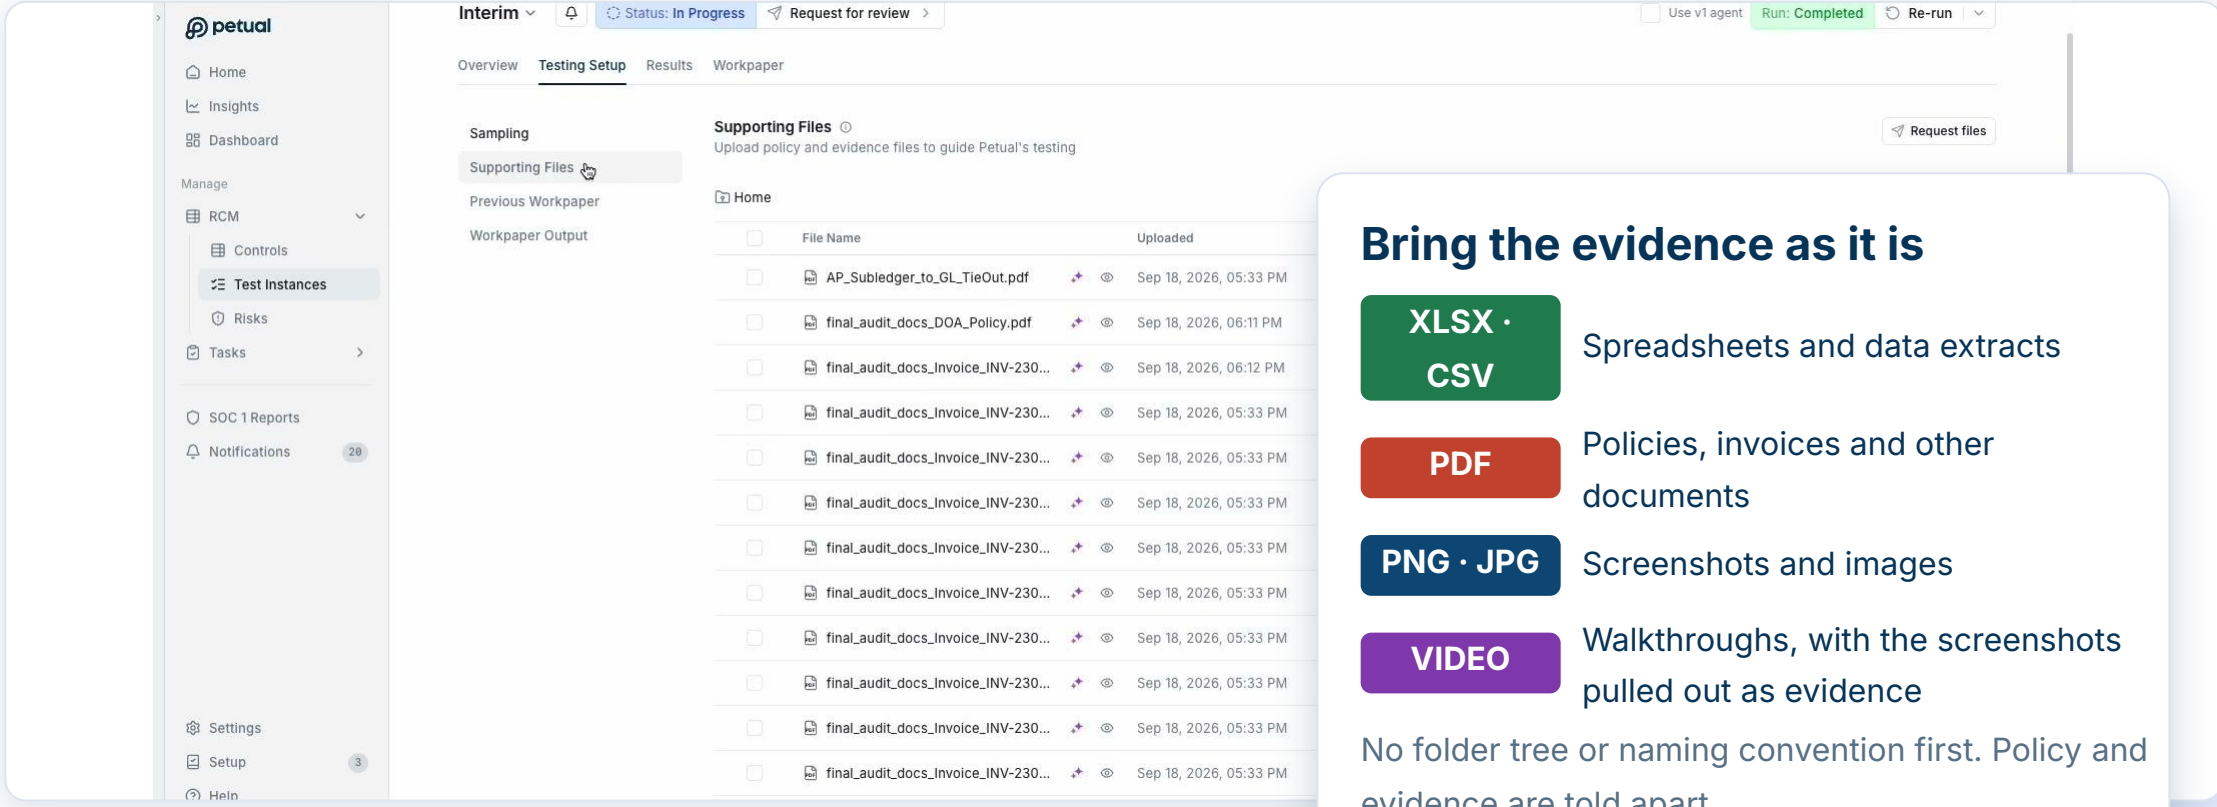


petual

1 Control | 2 Attributes | **3 Evidence** | 4 Samples | 5 Conclusion | 6 Workpaper

Harness at work · Tools and data
Reads each format; policy kept separate.

Drop the evidence in as it arrives



Bring the evidence as it is

XLSX · CSV — Spreadsheets and data extracts

PDF — Policies, invoices and other documents

PNG · JPG — Screenshots and images

VIDEO — Walkthroughs, with the screenshots pulled out as evidence

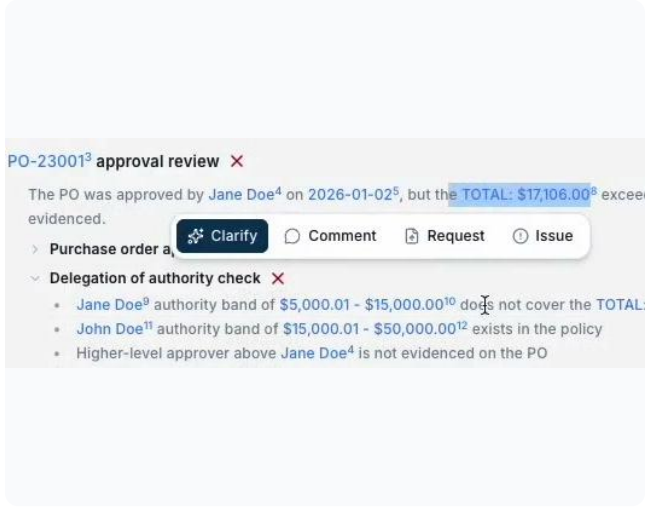
No folder tree or naming convention first. Policy and evidence are told apart.

petual

Harness at work · Memory
Clarifications carry into the next run.

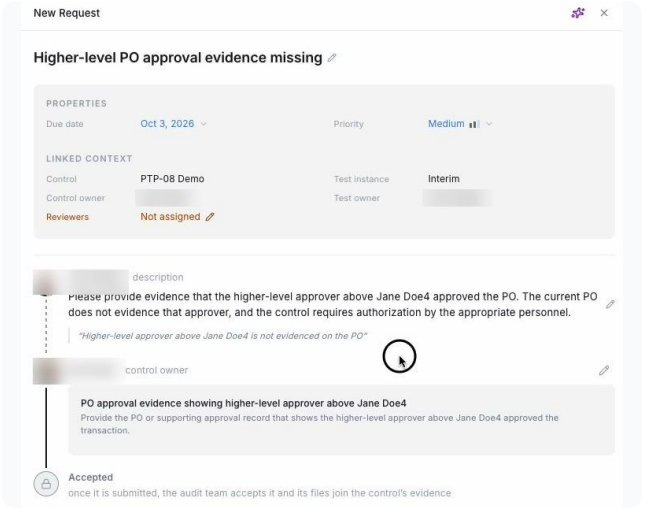
Missing something? Request it from the control owner

1 Flag the line that needs more



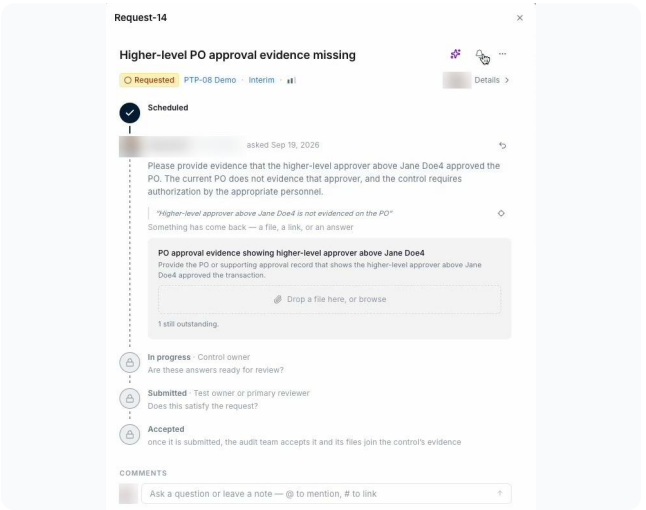
From any line of the reasoning: clarify it, comment on it, raise an issue, or request evidence.

2 The request carries the finding with it



Linked to the control, the test and the exact line that failed.

3 The owner answers in one place



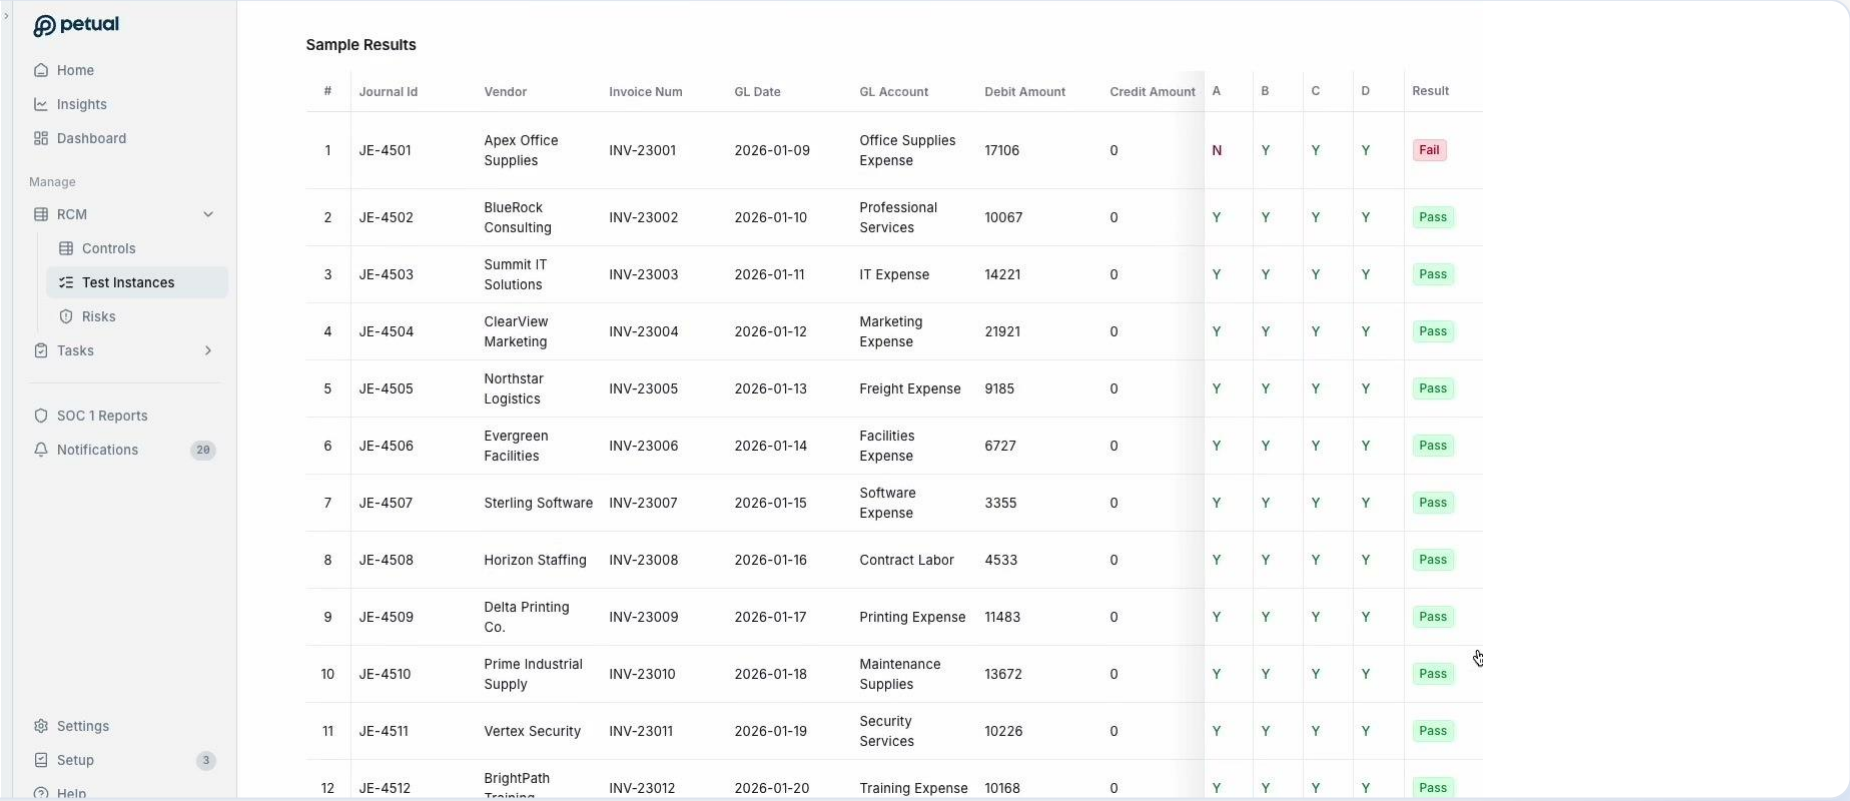
Tracked from requested to accepted. Accepted files join the control's evidence.

petual

1 Control | 2 Attributes | 3 Evidence | **4 Samples** | 5 Conclusion | 6 Workpaper

Harness at work · Orchestration
Every attribute tested on every sample. Most importantly, applied consistently

Every attribute, tested on every sample



Sample Results

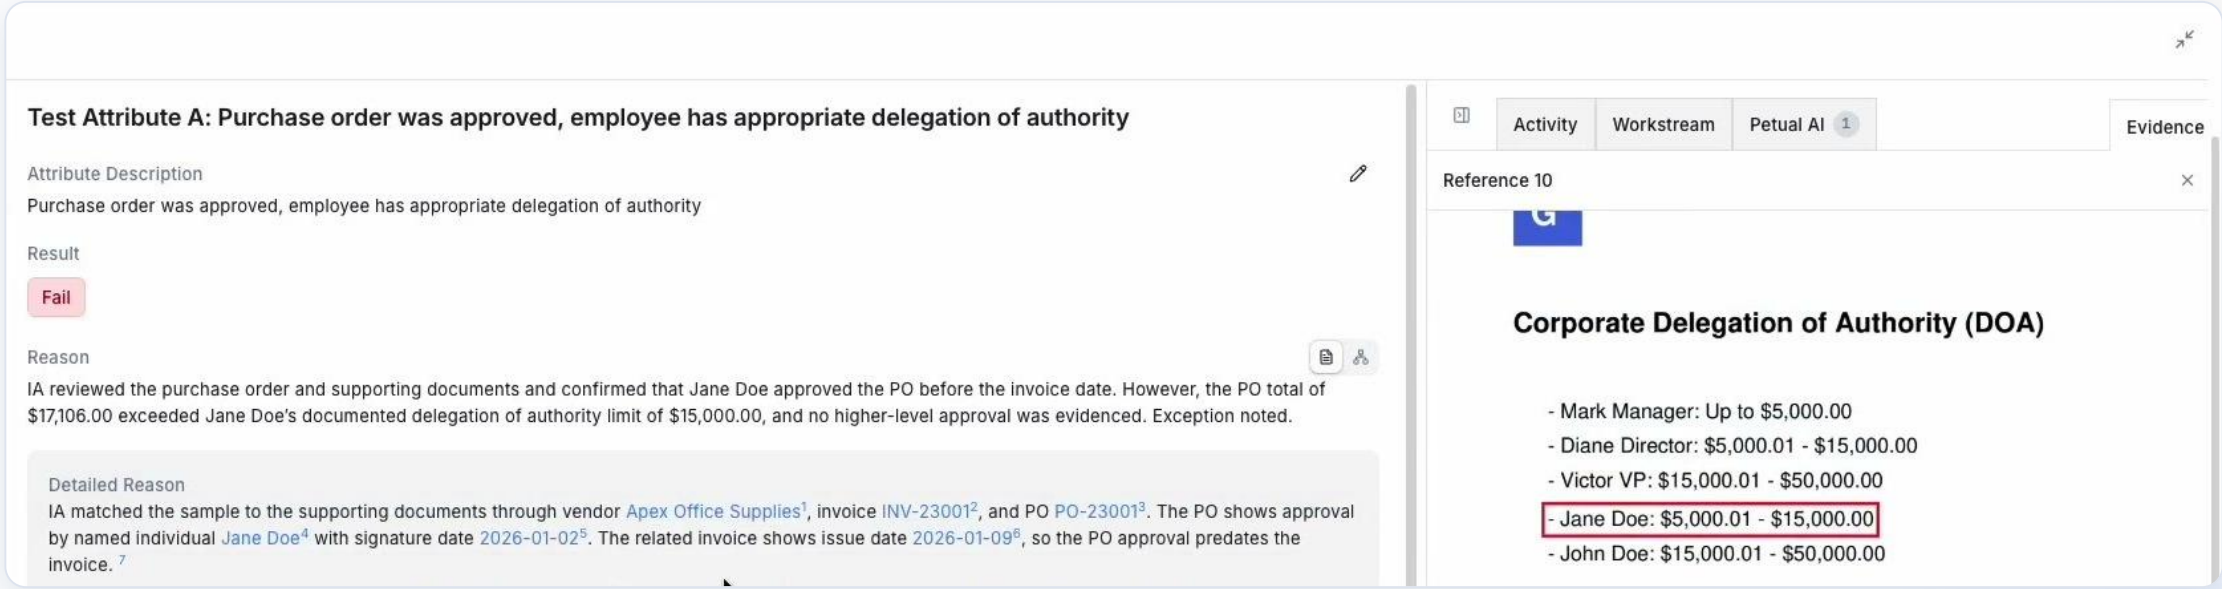
#	Journal Id	Vendor	Invoice Num	GL Date	GL Account	Debit Amount	Credit Amount	A	B	C	D	Result
1	JE-4501	Apex Office Supplies	INV-23001	2026-01-09	Office Supplies Expense	17106	0	N	Y	Y	Y	Fail
2	JE-4502	BlueRock Consulting	INV-23002	2026-01-10	Professional Services	10067	0	Y	Y	Y	Y	Pass
3	JE-4503	Summit IT Solutions	INV-23003	2026-01-11	IT Expense	14221	0	Y	Y	Y	Y	Pass
4	JE-4504	ClearView Marketing	INV-23004	2026-01-12	Marketing Expense	21921	0	Y	Y	Y	Y	Pass
5	JE-4505	Northstar Logistics	INV-23005	2026-01-13	Freight Expense	9185	0	Y	Y	Y	Y	Pass
6	JE-4506	Evergreen Facilities	INV-23006	2026-01-14	Facilities Expense	6727	0	Y	Y	Y	Y	Pass
7	JE-4507	Sterling Software	INV-23007	2026-01-15	Software Expense	3355	0	Y	Y	Y	Y	Pass
8	JE-4508	Horizon Staffing	INV-23008	2026-01-16	Contract Labor	4533	0	Y	Y	Y	Y	Pass
9	JE-4509	Delta Printing Co.	INV-23009	2026-01-17	Printing Expense	11483	0	Y	Y	Y	Y	Pass
10	JE-4510	Prime Industrial Supply	INV-23010	2026-01-18	Maintenance Supplies	13672	0	Y	Y	Y	Y	Pass
11	JE-4511	Vertex Security	INV-23011	2026-01-19	Security Services	10226	0	Y	Y	Y	Y	Pass
12	JE-4512	BrightPath Training	INV-23012	2026-01-20	Training Expense	10168	0	Y	Y	Y	Y	Pass

petual

1 Control | 2 Attributes | 3 Evidence | 4 Samples | **5 Conclusion** | 6 Workpaper

Harness at work · Checks
Each reference checked against its source.

Every conclusion traced to the evidence it rests on



The finding: the PO total was above the approver's documented limit, and no higher approval was evidenced.

The trail: each numbered reference opens the source, here the delegation of authority policy, with the line highlighted.

petual

Harness at work · Guardrails and logs
Every tickmark tied back to its source.

The workpaper: each tickmark tied back to the source

Tickmark	Description
A1	PO approval evidence names Jane Doe and shows signature date 2026-01-02.
A2	PO total $17,106.00 exceeds Jane Doe's documented DOA limit of $15,000.00; no documented temporary authority adjustment was provided — Missing Evidence
A3	PO approval date 2026-01-02 is earlier than invoice issue date 2026-01-09.
B1	Invoice unit price $17,106.00 agrees to PO unit price $17,106.00 for Office Supplies Expense; IA re-performed the comparison and noted 0.0% variance, within the 1% tolerance — Pass
B2	Invoice quantity 1 agrees to PO quantity 1 for Office Supplies Expense; IA re-performed the comparison and noted 0.0% variance, within the 1% tolerance — Pass
B3	Receipt quantity received 1 supports the invoiced quantity 1 for Office Supplies Expense — Pass
C1	PO approver Jane Doe and receipt recipient Alex Warehouse are different individuals, supporting segregation of duties — Pass
C2	PO number PO-23001 on the approved purchase order matches PO-23001 on the receipt — Pass
D	Booked debit amount $17,106.00 on JE-4501 matches invoice INV-23001 total due $17,106.00 — Pass
Source	final_audit_docs_PO_INV-23001.pdf
Location	page 1

G **GC Purchase Order**

PURCHASE ORDER
P.O. #: PO-23001 C2
Date: 2026-01-02

VENDOR:
Apex Office Supplies

SHIP TO:
Global Corp Industries
789 Corporate Way, Suite 100, Austin, TX 78701

Item Description	Qty	Price	Total
Office Supplies Expense	1 B2	$17,106.00	$17,106.00

A2, B1

TOTAL: $17,106.00

petual