

PORTFOLIO SAMPLE

AI Audit Automation in Regulated Financial Institutions

Interview-Led Long-Form B2B Editorial

Client: Adastra Group (enterprise AI & cloud platform)

Produced via: Podium IQ / Experience Bold AI

Format: Long-form thought leadership article (~2,300 words)

Source method: Original SME interviews — Omer Masudi & Canhan Fidanoğlu
(August 2026)

Content includes:

- Editorial brief — documenting interview findings, audience, EEAT approach, guardrails, and content outline
- Published article — the finished editorial output

Nabila Husseni

nabilahusseni1706@gmail.com · (437) 989-6925 · nabilahusseni.com

AI Audit Automation: Article

Draft V1 · 26 August 2026 · Built from the 24 August 2026 subject-matter-expert interview (Omer Masudi and Canhan Fidanoglu)

Field	Value
Article Name	AI Audit Automation: How Regulated Institutions Move from Point-in-Time Audits to Continuous Assurance
Date	August 26, 2026
Version #	V1
Target Audience	Internal audit leaders, chief auditors and audit and control teams at banks and other regulated financial institutions, including tier-2 and mid-size institutions still tracking audit findings in spreadsheets. Compliance and risk leaders under growing regulatory pressure who need audit coverage that is continuous rather than point-in-time. Technology and data leaders who get pulled into audit cycles as process, system and data owners and lose delivery capacity every audit season. Secondary and inferred rather than stated in the source interview: risk and audit functions in the other sectors where Adastra holds industry expertise (pharmaceuticals, manufacturing, automotive), plus public-sector audit bodies.
Business Objective	Establish Adastra's authority on audit automation ahead of product availability, using operator credibility (a product lead who ran transformation programs inside a bank and was audited by it; a builder with ten years in audit and control departments) rather than client proof, which does not yet exist. Give an internal audit or compliance leader a working mental model of what gets automated, in what order, and where the approach deliberately stops, so that a scoping conversation starts from an informed position. Answer the trust objection (AI hallucination versus defensible audit findings) directly, since it is the objection that otherwise ends the conversation.
Target Keywords	This article was written from a subject-matter-expert interview. Keywords: ai audit automation; audit automation software; continuous auditing; internal audit automation; ITGC audit automation; agentic AI for internal audit; sampling risk audit. No search volumes have been verified.
Target Queries	“what is AI audit automation”; “how to automate internal audit”; “continuous auditing vs periodic audit”; “does AI audit automation replace fraud detection”; “how can auditors trust AI findings”; “what is sampling risk in auditing”.
Pre-Publication Sign-Off	1. Omer to confirm his title and Canhan Fidanoglu's. 2. Decide the CTA destination. The article closes on a scoping conversation rather than availability, because the product is pre-launch. 3. Two items were deliberately left out for lack of verifiable sourcing and can be added if confirmed: the reported billion-dollar audit-automation

Field	Value
	investment by a Big Four firm. 4. No case study section. Adastra has no live implementation of this product; conversations are running in North America and Europe only.

SEO & GEO Metadata

Field	Value
Page URL	/insights/ai-audit-automation-continuous-assurance/ (suggested, confirm against site URL conventions)
Title Tag	AI Audit Automation: From Point-in-Time Audits to Continuous Assurance Adastra
Meta Description	AI audit automation turns regulations and policies into monitored rules, testing every transaction instead of a sample, with a human deciding every finding.
Canonical Tag	Self-referencing canonical to the final published URL above
Schema Tags	Article schema; FAQPage schema for the FAQ section (6 Q&A pairs); BreadcrumbList schema
Dates On Page	Display both a Published and an Updated date near the top. This matters more than usual here: the article describes a pre-launch product and will need revision once there is an implementation to cite.

Content

Component	Content
H1	AI Audit Automation: How Regulated Institutions Move From Point-in-Time Audits to Continuous Assurance
Body	AI audit automation ingests the regulations, standards and internal policies an organization is accountable to, works out which controls those documents require, monitors those controls against actual transactions and processes, and returns findings traced back to the source requirement. A person still decides what each finding means. That last sentence is the part most descriptions of the category leave out, and the part auditors care about most. What follows is the operating model, as Adastra is building it. Table of Contents • Why Audit Is Still a Manual, Point-in-Time Exercise • Where Audit Time Actually Goes • The Limit That Never Appears in the Audit Report • How AI Audit Automation Works, End to End • What It Catches: Two Worked Examples

Component	Content
	• What AI Audit Automation Is Not • How Auditors Can Trust an AI-Generated Finding • How to Measure the Business Case • Where the Audit Market Is Heading • What Makes Adastra's Approach Different • FAQs
H2	Why Audit Is Still a Manual, Point-in-Time Exercise
Body	In most regulated institutions the audit cycle is entirely manual and runs on a set frequency. Omer Masudi built Adastra's audit automation solution and previously ran large transformation programs inside a bank, describing each cycle as a fact-finding mission: auditors go out, collect information against the audit's requirements, then determine whether there are deficiencies. The cost is not confined to the audit function, because internal audit has to reach whoever owns the relevant processes, systems and data. Omer spent years on the receiving end, allocating five or six lead developers and architects to audit work and pulling them off the delivery they owned. There are two failure modes here, not one. The cycle is slow and expensive, which is the visible problem. The less visible one is that because it is point-in-time and depends on the right people surfacing the right information, there is no guarantee it reaches the right conclusion, and an incomplete finding is itself an exposure.
H2	Where Audit Time Actually Goes
Body	Here's how Omer breaks down where the hours actually go, based on running audits himself: 1. Understanding the requirements. Why an audit needs to run, when, and how often, and keeping current as regulations and internal standards change. This one stays a human responsibility after automation. 2. Planning. Identifying stakeholders, setting timelines, and sizing the work against how many audits are running and how complex each is. 3. Executing the audit. Chasing information, locating documentation, finding the right data across systems, and assessing it. 4. Managing findings. Producing them, building an action plan, and tracking last cycle's action items. Parts of stage three are already automated where the data sits in governed systems: an enterprise GRC platform holds the change record, so IT governance audits reach those conclusions faster. The work around the record is not. Someone still has to establish which artifacts a project produced, whether the right ones were approved by the right people at the right time, and where the money went. The information lives in systems. Assembling it is human labour.
H2	The Limit That Never Appears in the Audit Report
Body	A second constraint is structural rather than operational, and Canhan Fidanoğlu, who spent ten years in audit and control departments, puts it more sharply than most vendors would. A bank processes millions of transactions a day, so auditors sample. The illustration Canhan uses: if there are thousands of transactions and you can

Component	Content
	check a hundred, you are not covering almost everything with a small margin of error. You are missing 99% of them. That is a sampling risk, and it is one of the main risks any audit engagement carries. Assessing every transaction against the rule does not shrink it, it removes the reason it exists. That is the platform's design objective rather than a measured result.
H2	How AI Audit Automation Works, End to End
Body	The order of the stages matters. Most tooling in this space starts where the evidence is. This starts where the requirement is, with paper.
H3	From Documents to Monitorable Rules
Body	The first thing an institution loads is not data, it is documents: external regulations, a cybersecurity framework for instance, alongside internal standards and policies. That knowledge base is where every rule originates. From it, the platform works out what controls the institution needs and how each should be monitored. The sequence matters: requirements do not become rules directly, they become controls, and rules monitor the controls. A worked example. From a cybersecurity framework, a rule that every privileged account must have multi-factor authentication enabled, which reads operationally as: for each account in scope, check that MFA is present. The rule then watches the identity provider systems.
H3	The Citation Chain: Why an Auditor Can Defend the Finding
Body	This matters most to anyone who has had to explain a finding to a regulator. Every rule carries a citation back to the requirement that produced it. For the MFA rule, that is framework language requiring identities and credentials be issued, managed, verified, revoked and audited for all authorized devices, users and processes. The platform also flags where the same requirement appears in internal policy, so one violation breaches both. What the trace buys is a finding an auditor can defend without reconstructing how it was reached.
H3	Connectors, Agents and Continuous Monitoring
Body	Connectors reach the systems holding the data an audit needs: identity management for an ITGC audit, ERP and financial systems for a financial audit, or a warehouse that already receives the data. Omer is direct that this integration work is real and requires appetite from the institution to open its systems up. Once done, audits run on demand or continuously, and behind each sits a set of AI agents: an orchestrator, an agent doing evidence collection, and an auditor agent that executes the rules, detects anomalies and drafts narratives. The agents divide the audit labour. None of them decide anything.
H3	Findings: Severity, Citation, Remediation
Body	A completed run produces findings grouped by severity, categorized however the institution categorizes them. For the MFA rule, a finding reads: a user account was created without multi-factor authentication enabled. Each carries the citation of where the requirement comes from, the policies being violated, and a recommended remediation. When the remediation evidence is submitted, the control re-runs and a clean result closes the finding automatically. Opening tickets in the institution's own ticketing system remains on the roadmap. Administrative views cover the models and prompts in use, users and roles, and an audit log, because a function whose own controls are under audit will want the platform to be auditable too.

Component	Content
H2	What It Catches: Two Worked Examples
H3	An IT General Controls Exception
Body	A new user is granted privileged access, where the institution's standards require that access to be approved or to carry step-up authentication. The platform detects the account without it and raises the exception, because as written, the organization is offside its own standard. The difference from the status quo is timing: this surfaces at creation, not whenever a cycle happens to examine that sample.
H3	A Regulatory Reporting Exception
Body	In Canada, a transaction of \$10,000 or more conducted at a banking centre must be reported. A rule watching the transaction data finds one that went through with no corresponding report filed. Set the outcomes side by side: one flagged transaction raised and addressed, or the discovery at year end that 20,000 transactions lacked the required approval or report, followed by a self-declaration to the regulator.
H2	What AI Audit Automation Is Not
Body	It is not a fraud detection or transaction monitoring system and does not replace one. Institutions already have those, so this does not touch the customer experience. What it changes is exposure: the reputational and financial-loss risk arising from audit deficiencies gets contained proactively, rather than surfacing when an external audit finds it.
H2	How Auditors Can Trust an AI-Generated Finding
Body	Hallucination is a reasonable objection. Here it is addressed in four layers: 1. Design-level controls. Guardrails and defined skills constrain what the agents do, with challenger models where one model's output is verified by another. 2. Implementation and testing. Running against smaller subsets and reviewing the results, to establish the output is accurate and complete before it is relied on. 3. Human in the loop. Omer calls this the most important component. The platform is not built to take action; it presents findings and recommendations to a person, who decides and can reject a recommendation outright. 4. Ongoing assurance. Running samples in a lower environment periodically, to confirm the quality holds and no drift has crept in. AI projects stall, in Omer's account, when organizations reach a proof of concept and only then hit the governance and privacy questions they had not worked through. Designing for those at the foundation is what makes production viable.
H2	How to Measure the Business Case
Body	Asked about cost, Omer is direct that it depends on how manual the institution is today and how much of its audit estate it automates. Read the figures below with those conditions attached. • Audit-function time recovered. A ten-FTE audit function spends 40 to 50% of its time on keeping up to date, planning audits, running them and managing findings. The last three are what get cut down.

Component	Content
	• Manual effort automated. Almost 70 to 80% of total manual effort and process, with a human retained on every critical decision. • Cost reduction on automated cycles. 20 to 60%. Below 20% suggests something is not being done right; 60% is a good scenario. Savings scale with scope: two of fifteen cycles returns two cycles' worth. • Audit coverage. The full population rather than a sample, which is a change in kind rather than a percentage improvement. One anecdote is worth more than the arithmetic. In a recent conversation with the chief auditor of a tier-2 bank, Omer found an institution with no system for tracking audits at all. Findings were managed in a spreadsheet.
H2	Where the Audit Market Is Heading
Body	There is no mandate to automate audit today and Omer is not certain there ever will be, which makes it easier to defer and harder to defend deferring. What he expects instead: this becomes table stakes, on a timeline he will not put a number on. The change worth watching is the expectation rather than the technology. Audits that run on a predetermined frequency today could be expected to run continuously tomorrow, large incumbent audit firms are already committing serious capital on that premise, and regulators are clamping down on organizations of every size.
H2	What Makes Adastra's Approach Different
Body	The architectural claim comes first, because it is the checkable one. The platform is not built for a specific audit type. It is modular, and the requirement builds the audit process rather than the audit process being pre-built into the platform. An institution gets running by feeding in its own policies, standards and applicable regulations and building out the rules during implementation. Canhan draws the category contrast more precisely: most GRC solutions concentrate on the post-audit stage, finding follow-up and consolidation. This solution covers that too, but starts at the requirement. Underneath it is a data, AI and cloud company with expertise spanning financial services, pharmaceuticals, manufacturing and automotive.
H3	Deployment Options and What Comes Next
Body	Both options are going to be available: deployment into the institution's own cloud infrastructure, or run as a SaaS offering, on AWS or Azure, with tenancy on either side. For a function whose own controls are under audit, that choice is often the difference between a viable procurement and a non-starter. Canhan adds a forward-looking note: because the platform interprets objective source content and generates requirements from it, public-sector audit is a natural extension. A direction, not an offering today.
H2	FAQs
H3	What is AI audit automation?
Body	AI audit automation ingests the regulations, standards and internal policies an organization is accountable to, derives the controls those documents require, monitors them against actual transactions and processes, and returns findings traced to the source requirement. A person reviews every finding.

Component	Content
H3	Which parts of the audit process can actually be automated?
Body	Planning, evidence collection, rule execution and the drafting of findings. Keeping current with changing regulations stays a human responsibility, as does every decision about what a finding means and what action follows.
H3	Does AI audit automation replace fraud detection or transaction monitoring?
Body	No. Financial institutions already run fraud detection and transaction monitoring systems, and this neither overlaps with nor replaces them. It addresses control assurance and audit deficiencies, so it carries no direct customer-experience impact.
H3	How much of the manual audit effort can be automated?
Body	Almost 70 to 80% of total manual effort, with a human retained on every critical decision. Cost reduction on the cycles actually automated runs roughly 20 to 60%, and scales with the number of cycles in scope.
H3	How do you stop an AI system from hallucinating an audit finding?
Body	Four layers: guardrails and challenger models at the design level, testing against smaller subsets during implementation, a human in the loop who can reject any recommendation, and recurring sample runs in a lower environment to confirm no drift. Every finding also carries a citation to its source requirement.
H3	Can the platform be deployed in our own cloud environment?
Body	Yes, that is one of two intended options: deployable into the institution's own cloud infrastructure on AWS or Azure, or run as a SaaS offering. The product is pre-launch, so confirm current options before planning against them.
H2	Conclusion
Body	The shift here is narrower than "AI for audit." It is a move from audits that run on a frequency against a sample, to rules derived from the institution's own requirements and monitored continuously against all of it, with every finding traced back to the document that required it. The boundary is what makes the rest credible: the platform does the assembly and the assessment. A person still decides what every finding means.

Brief: AI Audit Automation in Regulated Financial Institutions

(AI Audit Automation. Aligned to article V1, 26 August 2026.)

Topic Selection Background:

AI Audit Automation in Regulated Financial Institutions

Why we chose this topic:

The topic comes from a subject-matter-expert interview on 24 August 2026 with Omer, who is building Adastra's audit automation product and previously ran large transformation programs inside a bank, and Canhan Fidanoğlu, who spent ten years in audit and control departments. Lexi Clarke was present but did not contribute substantive content. The interview establishes a problem Adastra can speak to first-hand: in most regulated organisations the audit cycle is still manual and point-in-time, and the cost lands on far more people than the audit function. Priority is high because the product is pre-launch, which means the article does credibility work rather than conversion work, and which sharply limits what it can claim.

What this topic will do about it:

The article explains what AI audit automation actually automates, stage by stage, and where it deliberately stops. It walks from a regulation or internal policy through to a monitored rule, a citation chain, a finding and a recommended remediation, using the two anomaly examples the interview provides. It then answers the objection that would otherwise end the conversation, that AI hallucinations and audit findings must be defensible, with the four layers of control Omer described. The target reader should finish understanding the operating model, not holding a capability list.

Content Brief, Section 1: Purpose & Best Practices

Topic: AI Audit Automation in Regulated Financial Institutions

Purpose and Goals

- Establish authority through lived operator experience rather than product marketing. Omer ran transformation programs inside a bank and was on the receiving end of its audits; Canhan spent a decade inside audit and control. That is the credibility the article trades on, because there is no client implementation to point to.
- Give an audit or compliance leader a clear model of what gets automated and in what order: requirements, planning, execution, findings management. Name the stages, then show which ones the platform touches.
- Make the sampling problem legible. Canhan's point is the sharpest and most citable thing in the interview: an auditor checking a hundred transactions out of thousands is missing 99% of them, not covering nearly all of them. Build a section around it.
- Answer the trust objection before it is raised. Auditors cannot act on an unexplainable finding, so the citation chain and the human-in-the-loop position are not features to list. They are the argument.

- Draw the product boundary explicitly, and do all of the above without implying the product is deployed. Stating plainly that this is not fraud detection builds more credibility than leaving it ambiguous.

Title: AI Audit Automation: How Regulated Institutions Move From Point-in-Time Audits to Continuous Assurance

Source Note: One Primary Source

Every claim traces to the 24 August 2026 interview, conducted by Williams Oleije and Nabila Husseni. No secondary research, competitor analysis, keyword data or third-party sources were used. Two consequences:

- If a fact is not in the interview, it does not go in the article without separate sourcing and sign-off. That includes anything a general reader might expect: market sizing, regulatory citations, named competitors, adoption statistics.
- Where the interview is uncertain, the article should be too. Omer hedged on cost percentages, on when automation became table stakes, and on Adastra's proof-of-concept-to-production rate. Reproduce the hedge rather than resolving it.

Audience & Tone Considerations

- Write for someone who runs audits, not someone learning what an audit is. Reference internal audit and the first, second and third lines as shared context and move on.
- The reader's scepticism is reasonable: they have been told before that a tool would fix this. Concede what is actually hard. Omer was open that the integration work is real and needs an appetite to open up source systems.
- Match the interview's register: operator to operator, specific, unhurried. Omer's most persuasive material is an anecdote with a number attached, such as the chief auditor tracking findings in a spreadsheet and the five or six lead developers pulled off delivery. Use those over abstractions.
- Favour audit-ready, defensible, traceable, continuous and proactive over revolutionary or game-changing. Keep the human in frame; the tone should not suggest auditors are being replaced, because Omer's position is the opposite.

EEAT & AI Citation Best Practices

- Open the article body with a self-contained 40 to 60 word definition, carrying the same definition as the AI Summary Snippet above it. The snippet block runs 100 to 150 words because it also covers the contrast and the transition.
- Attribute expertise by background, not job title. Neither speaker's title was stated on the call. Omer's bank operator history and Canhan's ten years in audit and control are the credentials that carry weight, and they are the ones the interview supports.
- Structure every H2 to stand alone as the answer to a question. The four-stage breakdown and the four trust layers should both be numbered lists; the measurement guidance should be a table.

- Use the two anomaly examples verbatim in substance: a privileged account created without step-up authentication, and a reportable transaction with no report filed. Specific, checkable examples get quoted; generic anomaly detection claims do not.
- Quote the citation-chain mechanism precisely. A rule traced to the requirement text in the source document, flagged where it also appears in internal policy, so one violation breaches both. That mechanism is the E-E-A-T story.
- Keep the six FAQ answers to 40 to 60 words each and independently liftable. Show a published and an updated date; the article will need revision once there is an implementation to cite.

Do Not Claim: Constraints and Guardrails

Two of these came from Omer directly when asked what to avoid: the fraud and financial crime exclusion, and the correction on cloud naming. The rest are this brief reading of his hedges. One caution on the source: asked whether there were claims to avoid, his answer was transcribed as “I don't think anything I said today is all true. As long as it's within those boundaries, I think we should be okay.” The sentence is garbled and the plausible reading is the opposite of the literal one, so confirm it against the recording. All eight are hard constraints on the draft.

1. No client implementations, named customers or case studies.

There are conversations with clients in North America and Europe, but no implementation anywhere and no real deployment. The product is not, in Omer's words, fully ready for prime time. Avoid softer phrasings like “institutions using the platform” too. Article V1 has no case study section for this reason.

2. No fraud detection, transaction monitoring or financial crime positioning.

Ruled out twice, and the only guardrail Omer volunteered unprompted. Banks already have those systems and this does not replace them. One unresolved tension: the \$10,000 reportable-transaction example he gave is itself an anti-money-laundering obligation, so it sits inside the category he asked to avoid. V1 keeps it as a regulatory reporting control, which is an editorial judgement, not his. Clear it with him and have a second IT general controls example ready.

3. No customer-experience or false-positive claims.

Asked about false positives, delayed investigations and customer friction, Omer's answer was that audit automation creates no direct client-experience impact. Do not import benefits from the fraud-detection category.

4. The interface shown was a prototype, and its numbers are illustrative.

Omer used a simplified clickable prototype rather than the actual application. Every figure on those screens, the roughly 4,000 auditor hours saved and the counts of critical, high, medium and low findings, is demo data showing what the dashboard reports. V1 describes the hours-saved tile as a reported metric and omits both figures. Any screenshot must be labelled a product concept.

5. Name AWS and Azure only.

Omer began to say the platform runs on any cloud, corrected himself mid-sentence, and settled on AWS and Azure for now. Do not name Google Cloud or GCP.

6. Publish cost impact as a conditional range, never a single number.

Omer's answer was openly a consulting answer: it depends on how manual the organisation is, what systems it runs, and how much it automates. His range was 20 to 60%, with below 20% suggesting something was not being done right, and savings proportional to how many cycles are automated. He allowed 70% in passing. Publish 20 to 60% with the conditionality attached; do not lead with 70%, and do not present any figure as a benchmark.

7. Attribute framing to the speaker, not to the industry.

The four-stage model is Omer's own read; he said plainly it is not an industry methodology. He reached for "one size fits all," flagged that he does not use the term, then explained what he meant: the platform is modular and the requirement builds the audit process. Use the explanation, not the shorthand. Keep his competitive hedge intact too, since "I have not seen this approach" is not "no one else does this."

8. Verify or omit the three unsourced items.

The reported billion-dollar audit-automation investment by a Big Four firm and the enforcement actions Omer named were recalled in conversation, and the enforcement examples sit on the financial crime side the article avoids. The proof-of-concept-to-production rate, recalled as roughly 75%, was hedged and unconfirmed by anyone present. V1 omits all three: the market section says incumbents are committing serious capital without the figure, and the trust section makes the production argument qualitatively. Add any of them back only once sourced.

TARGET AUDIENCE	Internal audit leaders, chief auditors and audit and control teams at banks and other regulated financial institutions, including tier-2 and mid-size institutions still tracking findings in spreadsheets. Compliance and risk leaders who need audit coverage that is continuous rather than point-in-time. Technology and data leaders who get pulled into audit cycles as process, system and data owners.
CUSTOMER STAGE	Problem-aware to solution-aware. Deliberately not bottom-of-funnel: the product is pre-launch with no live implementations, so the article educates and builds credibility rather than driving a purchase decision.
PAGE TITLE	AI Audit Automation: From Point-in-Time Audits to Continuous Assurance Adastra
META DESCRIPTION	AI audit automation turns regulations and policies into monitored rules, testing every transaction instead of a sample, with a human deciding every finding.
H1	AI Audit Automation: How Regulated Institutions Move From Point-in-Time Audits to Continuous Assurance
WORD LENGTH	2,000 to 2,300 words. V1 as drafted runs 2,314 including headings. It runs longer than a standard explainer because of the four-stage breakdown, the end-to-end architecture walkthrough, two worked examples, a four-layer trust model and six FAQs. To reach a hard 2,000, cut the administrative-views sentence in Findings, halve the market section, and take roughly 15 words from each FAQ answer. Do not cut the citation chain or the trust layers.
KEYWORDS & CROSS-LINKING	Not set. No keyword research or site audit was in scope for this brief. Seed terms and suggested queries are carried in the article document's metadata table for an SEO pass to validate before publication.

Alignment With Article V1, and Open Items

This brief was reconciled against the drafted article on 26 August 2026. Four things changed in the article that this outline now reflects, so a reviewer comparing the two documents should find no structural gaps.

- The end-to-end section was consolidated from eight subsections to four: From Documents to Monitorable Rules, The Citation Chain, Connectors and Agents and Continuous Monitoring, and Findings. Knowledge and rules were merged, connectors and agents and monitoring were merged, and the dashboard walkthrough folded into Findings. Nothing was dropped, only combined, and the citation chain kept its own subsection because it does the heaviest work.
- Deployment Flexibility and the public-sector extension became one H3, Deployment Options and What Comes Next, because both are short and both are forward-looking.
- The measurement table went from six metrics to four. Speed of detection and residual exposure were cut as metrics because the article already makes both points in prose, in the regulatory reporting example and in What AI Audit Automation Is Not.
- A manual-versus-continuous comparison table appeared in an earlier version of this brief and is not in the article. It was cut rather than kept as an unbuilt instruction, because every row restated a point the surrounding prose already makes. If a visual asset is wanted later, that table is the obvious candidate.

Open items before publication

1. Confirm Omer's title and Canhan Fidanoğlu's. The article currently describes both by background instead.
2. Get Omer's approval on the \$10,000 reportable-transaction example (see Do Not Claim, point 2), and have a second IT general controls example ready if he declines it.
3. Verify or drop the three unsourced items (see Do Not Claim, point 8). All three are currently out of the article.
4. Decide the CTA destination. The article closes on a scoping conversation rather than availability, because the product is pre-launch.

Content Brief, Section 2: On Page Content Outline

The outline below matches article V1 exactly, H2 for H2 and H3 for H3.

AI Summary Snippet

Format note: 100 to 150 words at the top of the article. Definition first, then straight into the contrast. No product name in this block.

- Define it once, plainly: the platform ingests the regulations, standards and internal policies the organisation is accountable to, works out which controls those documents require, monitors them against actual transactions and processes, and returns findings traced back to the source requirement, with a person deciding what each finding means.
- Keep Omer's starting point visible: the chain begins with paper, regulator requirements and internal standards and policies, and runs soup to nuts from there.

- Then the contrast in one sentence, plus the pre-launch disclosure: most audits run on a fixed frequency against a sample, this runs continuously against all of it, and nothing in the article is a client result.

Table of contents: add below the snippet, consistent with prior briefs. Eleven H2 entries.

H2:

Why Audit Is Still a Manual, Point-in-Time Exercise

Format note: Short. The reader lives here already, so name the cost precisely and move on.

- The cycle is entirely manual and runs on a set frequency. Omer calls each one a fact-finding mission: auditors go out, collect information against the audit's requirements, then determine whether there are deficiencies.
- The cost is not confined to the audit function. Internal audit has to reach whoever owns the relevant processes, systems and data. Use Omer's own account of allocating five or six lead developers and architects to audit work and pulling them off delivery.
- Two failure modes, not one. Slow and expensive is the visible problem. The invisible one is that a point-in-time audit dependent on the right people surfacing the right information carries no guarantee of the right conclusion, and an incomplete finding is itself an exposure.

H2:

Where Audit Time Actually Goes

Format note: Numbered list, and the strongest candidate for AI Overview extraction. Introduce it as Omer's own read, not an industry standard; he was explicit about that.

1. Understanding the requirements. Why an audit runs, when, and how often, plus keeping current as regulations and internal standards change. Stays a human responsibility after automation.
 2. Planning. Identifying stakeholders, setting timelines, sizing the work against how many audits are running and how complex each is.
 3. Executing the audit. Chasing information, locating documentation, finding the right data across systems, and assessing it.
 4. Managing findings. Producing them, building an action plan, tracking last cycle's action items.
- Follow the list with the honest caveat about what is already automated. Enterprise GRC platforms hold the change record, so IT governance audits reach those conclusions faster. The work around the record is not automated: which artifacts a project produced, whether the right ones were approved by the right people at the right time, where the money went.
 - Land the framing: the information lives in systems, but assembling it is human labour.

H2:

The Limit That Never Appears in the Audit Report

Format note: Short, self-contained and the most quotable section. Attribute to Canhan Fidanoğlu and lead with the ten years in audit and control. The interview does not establish pronouns, so use the name.

- A bank processes millions of transactions a day, so auditors sample. State it as Canhan did: if there are thousands of transactions and you can check a hundred, you are missing 99% of them. That is sampling risk, one of the main risks any audit engagement carries.

- Keep the figures as the illustration they were. Canhan said “thousands,” not a specific number, so do not tighten it into arithmetic that has to hold up.
- Assessing every transaction against the rule removes the reason sampling risk exists rather than shrinking it. Write it as the platform's design objective, not a measured result.

H2:

How AI Audit Automation Works, End to End

Format note: The structural centre of the article. Follow the data flow, requirement to rule to evidence to finding, rather than listing modules. Four H3s. Any screenshot must be labelled a product concept.

- Lead-in, two sentences: most tooling in this space starts where the evidence is, this starts where the requirement is, with paper.

H3:

From Documents to Monitorable Rules

- The first thing loaded is documents, not data: external regulations such as a cybersecurity framework, alongside internal standards and policies. That knowledge base is where every rule originates.
- The step not to skip, because the dashboard reports on it later: the platform works out what controls the institution needs and how each should be monitored. Requirements do not become rules directly. They become controls, and rules monitor the controls.
- Worked example: from a cybersecurity framework, a rule that every privileged account must have multi-factor authentication enabled, expressed operationally as check that MFA is present for each account in scope. The rule then watches the identity provider systems.

H3:

The Citation Chain: Why an Auditor Can Defend the Finding

Format note: Its own H3, not folded into rules. This is the E-E-A-T centre and the direct answer to the trust objection later.

- Every rule carries a citation back to the requirement that produced it. In the interview's example, framework language requires that identities and credentials be issued, managed, verified, revoked and audited for all authorised devices, users and processes.
- The same requirement is flagged where it also appears in internal policy, so one violation breaches both the internal standard and the regulator's requirement. Every rule and every finding carries the trace.
- Say what that buys: a finding an auditor can defend without reconstructing how it was reached.

H3:

Connectors, Agents and Continuous Monitoring

- Connectors reach the systems holding the data an audit needs: identity management for an ITGC audit, ERP and financial systems for a financial audit, or a warehouse that already receives the data.
- Do not soften the honest part. Omer was direct that the integration work is real, needs investment, and needs appetite to open the institution's systems up. This is the paragraph a technology leader reads most closely.
- Behind each audit run sits a set of agents. Keep the division as stated: an orchestrator, an agent doing evidence collection, and an auditor agent that executes rules, detects anomalies and drafts narratives. Three roles, not four or five.
- Make the point explicit: the agents divide the audit labour, and none of them decide anything. Some rules are not invoked at a phase in a cycle at all but run continuously, which is what changes the cadence.

H3:

Findings: Severity, Citation, Remediation

- A completed run produces findings grouped by severity, categorised however the institution categorises them. Use the MFA example: an account created without multi-factor authentication enabled.
- Each finding carries the citation of where the requirement comes from, the policies being violated, and a recommended remediation. Label as roadmap: opening tickets in the institution's own ticketing system, and monitoring findings to closure.
- Close on the administrative views, which models and prompts are in use, users and roles, an audit log. One sentence, but include it: a function whose own controls are under audit will want the platform to be auditable too.

H2:

What It Catches: Two Worked Examples

Format note: Both came straight from the interview. Write each as a short narrative covering the rule, the exception and what happens next, rather than a bullet. These are the passages most likely to be quoted.

H3:

An IT General Controls Exception

- A new user is granted privileged access where standards require approval or step-up authentication. The platform detects the account without it and raises the exception, because as written the organisation is offside its own standard.
- The contrast: this surfaces at creation, not whenever a cycle happens to examine that sample.

H3:

A Regulatory Reporting Exception

- In Canada a transaction of \$10,000 or more at a banking centre must be reported. A rule watching the transaction data finds one that went through with no report filed, and asks what happened.

- Then the comparison that carries the section: one flagged transaction raised and addressed, versus discovering at year end that 20,000 transactions lacked the required approval or report and self-declaring to the regulator.
- Write it as a regulatory reporting control. It is not fraud detection, and it needs the sign-off in Do Not Claim, point 2.

H2:

What AI Audit Automation Is Not

Format note: A short paragraph. Counter-intuitively a trust-building section, and also a hard constraint. The boundary is not optional.

- Not a fraud detection or transaction monitoring system, and it does not replace one. Institutions already have those.
- It therefore does not touch the customer experience, so no claims about false positives or investigation delays.
- What it changes is exposure: reputational and financial-loss risk from audit deficiencies contained proactively, rather than surfacing when an external audit finds it.

H2:

How Auditors Can Trust an AI-Generated Finding

Format note: The objection-handling section, and it should read as an answer rather than a reassurance. Four numbered layers in Omer's order, with the human-in-the-loop layer given the most weight since he called it the most important component.

1. Design-level controls. Guardrails and defined skills constrain the agents, plus challenger models where one model's output is verified by another.
 2. Implementation and testing. Running against smaller subsets and reviewing results, to establish the output is accurate and complete before it is relied on.
 3. Human in the loop. The one to lead on. The platform is not built to take action; it presents findings and recommendations to a person, who decides and can reject a recommendation outright.
 4. Ongoing assurance. Samples run in a lower environment periodically, to confirm quality holds and no drift has crept in.
- Close on why this matters commercially. Omer's argument: AI projects stall when organisations reach a proof of concept and only then hit the governance and privacy questions they had not worked through. Designing for those at the foundation is what makes production viable.
 - If the proof-of-concept-to-production rate is ever confirmed, it belongs here with attribution. V1 makes the point qualitatively (see Do Not Claim, point 8).

H2:

How to Measure the Business Case

Format note: Table, then two or three sentences of narrative. Every figure needs its condition attached, because a bare percentage will be quoted back without it.

Metric in the article	What the interview supports
Audit-function time recovered	A ten-FTE audit function spends 40 to 50% of its time on keeping up to date, planning audits, running them and managing findings. The last three are what get cut down; keeping current stays a human responsibility.
Manual effort automated	Almost 70 to 80% of total manual effort and process, with a human in the loop retained for every critical decision.
Cost reduction on automated cycles	20 to 60%. Below 20% suggests something is not being done right; 60% is a good scenario. Proportional to scope: two of fifteen cycles returns two cycles' worth.
Audit coverage	The full population instead of a sample. A change in kind rather than a percentage improvement, and the payoff of the sampling-risk section above.

- Frame it as Omer did: the answer depends on how manual the institution is today and how much of its audit estate it automates, and the honest version of that answer is a range with its conditions attached.
- Close on the anecdote about the tier-2 bank whose chief auditor had no system for tracking audits at all, with findings managed in a spreadsheet. Omer's read: for a growing institution that is not a reason to wait.

H2:

Where the Audit Market Is Heading

Format note: Reproduce Omer's uncertainty rather than resolving it. He was careful here. The two third-party facts stay out unless sourced (see Do Not Claim, point 8).

- No mandate exists to automate audit today and Omer is not sure one ever will, which makes it easier to defer and harder to defend deferring.
- What he expects instead: this becomes table stakes on a timeline he will not put a number on, from a few years to a decade.
- The structural shift worth naming: the change is in the expectation, not the technology. Audits on a predetermined frequency today could be expected to run continuously tomorrow, incumbent audit firms are already committing serious capital on that premise, and regulators are clamping down on organisations of every size.

H2:

What Makes Adastra's Approach Different

Format note: Lead with the architectural claim, because it is the checkable one, not with the company credentials.

- The architectural claim in Omer's words rather than his shorthand: the platform is not built for a specific audit type. It is modular, and the requirement builds the audit process rather than the audit

process being pre-built into the platform. An institution gets running by feeding in its own policies, standards and applicable regulations and building out the rules during implementation. Avoid “one size fits all.”

- The category contrasts with Canhan: most GRC solutions concentrate on the post-audit stage of finding follow-up and consolidation, which this platform also covers. The difference is that it starts at the requirement.
- Company credentials, two sentences at most: a data, AI and cloud company with expertise spanning financial services, pharmaceuticals, manufacturing and automotive, and a product built with that cross-industry knowledge of how audits actually run.

H3:

Deployment Options and What Comes Next

Format note: One H3 covering both, in the tense Omer used. He was describing intent for a pre-launch product.

- Both options are going to be available: the institution's own cloud infrastructure, or a SaaS offering, on AWS or Azure, with tenancy on either side. Why it matters: a function whose own controls are under audit often cannot put regulated data in someone else's tenancy.
- Then Canhan's forward-looking note, clearly labelled as direction rather than an offering. Because the platform interprets objective source content, regulatory frameworks, legal frameworks and internal procedures, and generates requirements from it, public-sector audit is a natural extension. Public auditors do precisely this work.

H2:

FAQs

Format note: Six questions, 40 to 60 words each, independently liftable. Implementation-level, not definitional. Each question below is an H3 in the article.

H3:

What is AI audit automation?

- Restate the opening definition in a form that stands alone: documents in, controls derived, monitored against real transactions and processes, findings traced to the source requirement, a person deciding what happens next.

H3:

Which parts of the audit process can actually be automated?

- Planning, evidence collection, rule execution and the drafting of findings. Keeping current with changing regulations stays a human responsibility, as does every decision about what a finding means.

H3:

Does AI audit automation replace fraud detection or transaction monitoring?

- The answer is no, and the answer needs to be that blunt. See Do Not Claim, point 2. Close on the fact that there is therefore no customer-experience impact to claim.

H3:

How much of the manual audit effort can be automated?

- Almost 70 to 80% of total manual effort, with a human on every critical decision, and 20 to 60% cost reduction on the cycles actually automated. Carry the conditions; see Do Not Claim, point 6.

H3:

How do you stop an AI system from hallucinating an audit finding?

- Compress the four layers into one answer, then add the citation chain as the fifth reassurance: every finding traces back to its source requirement.

H3:

Can the platform be deployed in our own cloud environment?

- Yes, as one of two intended options, on AWS or Azure, with tenancy on either side. End on the pre-launch caveat: confirm current options before planning against them.

H2:

Conclusion

- One sentence restating the shift: from audits run on a frequency against a sample, to rules derived from the institution's own requirements and monitored continuously against all of it, with every finding traced to the document that required it.
- One sentence on the boundary, because it is what makes the rest credible: the platform does the assembly and the assessment, and a person still decides what every finding means.
- A single closing CTA framed as a conversation about the reader's own audit cycles. Not availability, not a trial, no implied deployments. Confirm the destination before publication.