

Specimen — illustrative example

Kestrel Ledger Inc. is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

AI risk advisory report

Kestrel Ledger Inc. · Written and reviewed by RiskWoven.com

Use case: In scope: extraction of invoice fields, matching to purchase orders and receipts, anomaly flagging, and the support copilot over our own knowledge base. Explicitly out of scope: any automatic release of payment, any credit or supplier scoring, any use of customer data to train models for other customers. The coding assistant and the meeting notetaker are internal tools and are not part of the product.

Score 81/100 · Critical · Confidence Low · Intake 100% complete

Executive summary

Kestrel Ledger Inc. recorded a 81/100 critical AI risk score for the described use case. 4 critical advisory conditions require customer action. The score uses ten weighted domains, disclosed contextual adjustments and evidence-gated control credit.

Conditions that need action

- Sensitive information sent to a provider that may retain or train on it. Establish in writing what the provider may do with prompts, inputs and outputs, and confirm by technical validation that the data actually sent matches what the contract permits. Where de-identification is relied on, test it against the identifiers that appear in real records rather than assuming the transformation is complete.

3 risk domains have not been assessed: Safety and harmful outcomes, Fairness and human rights, Operational resilience and continuity. Until each is assessed it is scored at its maximum, which is why the score is 81.

Analysis

Kestrel Ledger operates four AI systems in production: an invoice extraction and anomaly flagging model fine-tuned in house, a support copilot backed by a commercial model provider API, a third-party meeting notetaker used on customer sales calls, and a commercial coding assistant. Only the extraction model is part of the product sold to customers, but all four carry unmanaged risk. The most

consequential exposure is in the extraction model: it reads supplier bank account numbers, its accuracy has not been measured since 2022, drift monitoring does not exist, and a missed or wrong read could direct a payment to a fraudulent account. Mandatory human review before payment release is the primary safeguard, but reviewer behaviour is not measured and automation bias would be invisible until money moved to the wrong place. Three AI vendors are in production with no security assessment, no data processing agreement with the meeting notetaker, and no retained copy of the model provider terms. Customer invoice metadata is held in an AWS us-east-1 warehouse with no transfer assessment and no customer notice, and the lawful basis for using customer invoices to fine-tune a shared model has not been established. Three legal questions have been identified and none has been put to counsel. Governance is informal: no AI policy, no board oversight of AI risk, no completed impact assessment for any system, and no incident reporting process in operation. The conditions set by Dana Whitcombe, Chief Technology Officer, require a security and privacy assessment of the three unreviewed vendors, a data processing agreement or removal of the meeting notetaker, a contract position on fine-tuning, a transfer assessment for us-east-1, a first restore test, and a documented accuracy baseline, all targeted by end of Q1 2027. None of these conditions has an assigned owner today.

What drives the risk

- **Accuracy and reliability:** This is the domain most directly connected to Kestrel Ledger's stated worst-case outcome: a payment released to a fraudulent or altered supplier bank account because the extraction model read the wrong account number and the anomaly step did not flag it. The accuracy baseline has not been measured since 2022, two retraining cycles shipped without a documented threshold or formal confirmation, and drift monitoring does not exist. The inherent impact is scored at the highest level in the intake. Controls are reported but unvalidated, and the methodology credits only validated controls. This domain carries the greatest weight in the overall score.
- **Privacy and data protection:** Customer invoice metadata and extracted fields, including sole-proprietor banking details that are personal information under PIPEDA, are held in an AWS us-east-1 warehouse with no transfer assessment, no contractual safeguard and no customer notice. The lawful basis for fine-tuning on customer invoices has not been established. The meeting notetaker holds sales call recordings under unknown terms with no data processing agreement. No privacy impact assessment has been started for any of the four systems. The inherent impact and likelihood are both scored at four in the intake, giving this domain the highest inherent risk product of any assessed domain.
- **Third-party and vendor dependency:** Three AI vendors are in production with no security assessment, no data processing agreement with the meeting notetaker, and no retained copy of the model provider terms. The meeting notetaker joins customer sales calls and its storage location and retention period are unknown. The inherent impact and likelihood are both scored at four in the intake, matching privacy as the highest inherent risk product. Control effectiveness is assessed at ten percent in the intake, the lowest of any domain, because only AWS and Entra ID have been through any review.

- Legal and regulatory obligations: Three legal questions are now written down and none has been put to counsel. The fine-tuning use of customer invoices, the meeting notetaker recording practice, and the us-east-1 transfer are all knowing exposures. Control effectiveness is assessed at ten percent in the intake. The target date for legal consultation is December 2026, which is the earliest of any domain action, reflecting the urgency Kestrel itself has assigned.
- Cybersecurity and misuse: Privileged access to a production environment holding customer supplier banking data is unreviewed, leaver revocation is unevidenced, and CloudTrail logging is unread with no alerting on privileged actions. The 2023 penetration test findings were never tracked to closure. Three contractors use unmanaged personal devices. The inherent impact is scored at the highest level in the intake and control effectiveness at twenty percent, reflecting that the technical controls exist but produce no evidence of operation.

Watch items

- The accuracy of the extraction model on supplier bank account numbers is unknown. The 2022 baseline has not been re-measured and drift monitoring does not exist. Until a current baseline exists, every retraining cycle ships with unverified accuracy on the field most directly connected to financial loss.
- Customer invoices are being used today to fine-tune a model that serves other customers. The customer agreements do not clearly permit this. Counsel has not been consulted. This is a knowing exposure and should be stopped or contractually resolved before the next retraining cycle.
- The meeting notetaker joins customer sales calls with no data processing agreement, no written notice, and no known retention period or storage location. The compensating control of disabling it on customer-facing calls should be confirmed as implemented immediately.
- Invoice metadata and extracted fields, including sole-proprietor banking details, are held in AWS us-east-1 with no transfer assessment and no customer notice. The compensating control of stopping new data flowing into that warehouse should be confirmed as implemented.
- The 2023 penetration test produced findings that were never tracked to closure. No security testing has been done since. The environment holds customer supplier banking data.
- Three AI vendors — the model provider, the meeting notetaker and the coding assistant — have had no security assessment and no SOC 2 or equivalent has been requested. None of the conditions set by Dana Whitcombe, Chief Technology Officer, has an assigned owner today.
- Override, rejection and time-to-approve rates for customer payment reviewers are not measured. Automation bias — reviewers accepting AI output without independent checking — would be invisible until a payment went to the wrong account.
- The copilot has no tested fallback if the model provider API is withdrawn. A withdrawal would take the copilot down immediately with no recovery path.
- Known model limitations — weaker anomaly flagging for suppliers with fewer than six months of history, degraded extraction on handwritten and low-resolution invoices — are not disclosed in

customer-facing documentation. Customers cannot account for these limitations in their review decisions.

- No AI policy, no board oversight of AI risk, no AI literacy programme, no incident reporting process in operation, and no internal audit of AI exist. Almost every control in every domain is unverifiable because the governance structure that would produce and retain evidence has not been built.

Scope

Sector: Technology. Organisation size: 101 to 500 staff. Jurisdiction: Ontario. Uses: Writing, summaries and productivity, Customer service or chatbots, Analytics and business recommendations, Cybersecurity, fraud or safety decisions, Software development or technical operations. Information: Personal information, Financial or payment information, Confidential business information, Credentials, secrets or security information. Framework basis: ISO-based. Additional requirements for your review: The organisation has real security engineering but no management system. Nothing that exists is written down in a form a certification body or an enterprise customer can read, and the AI systems in particular have had no assessment of any kind. The most serious exposures are the untested restore, the three unassessed AI vendors, the undocumented cross-border transfer, and the absence of any measured accuracy baseline for a model that decides what a payment approver sees.

Risk register

Priority	Domain	Risk and why	Treatment	Owner	Target
Critical: act before continued use	Fairness and human rights	Error rates differ between groups of affected people, so some are flagged or missed more often than others. Bias and fairness testing has not been done. You noted: Nothing systematic. Exception queues are worked by our own operations team, which catches low-quality scans before they reach the customer, but that is a staffing arrangement rather than a control ...	Test error rates across the groups the system affects, record the results, and set a threshold that stops use if the gap is too wide.	Dana Whitcombe, Chief Technology Officer	Before continued use

Priority	Domain	Risk and why	Treatment	Owner	Target
Critical: act before continued use	Privacy and data protection	Personal information is collected and used without an established legal authority or consent. The lawful basis for processing is recorded as not established.	Establish and document the privacy authority for each purpose, including any secondary use, before further processing.	Dana Whitcombe, Chief Technology Officer	2027-03
High	Accuracy and reliability	The extraction model reads a supplier bank account number wrongly, or reads an altered one as legitimate, and the anomaly step does not flag it. Recorded by you in the domain assessment.	Re-establish a measured accuracy baseline against human keying on a held-out sample, reported monthly per customer and per invoice type. Set an acceptance threshold that blocks a model release. Add alerting on extraction confidence and on the rate of anomaly flags. Produce the approval log extract as standing evidence.	Marcus Oyelaran, Director of Platform Engineering	2027-03
High	Privacy and data protection	Customer invoice metadata and a sample of extracted fields sit in an analytics warehouse in AWS us-east-1 with no transfer assessment, no contractual safeguard and no notice to customers. Recorded by you in the domain assessment.	Complete a transfer assessment for us-east-1 and decide whether the warehouse moves to ca-central-1 or the transfer is documented and disclosed. Write the internal privacy standard that the public policy implies. Settle in the customer agreement whether invoices may be used to fine-tune a shared model, and stop doing so until that is settled.	Dana Whitcombe, Chief Technology Officer	2027-03

Priority	Domain	Risk and why	Treatment	Owner	Target
High	Cybersecurity and misuse	Privileged access to production is granted through Entra ID and nobody reviews it. Joiner and leaver handling is a ticket to IT with no evidence of timely revocation, access reviews happen when someone remembers, and logging is enabled but ... Recorded by you in the domain assessment.	Run and record a full privileged access review, then put it on a quarterly cycle. Put leaver revocation on a measurable service level with evidence retained. Turn on alerting for privileged actions and root account use. Bring contractor devices under management or restrict what they can reach. Resume penetration testing and track findings to closure.	Marcus Oyelaran, Director of Platform Engineering	2027-02
High	Transparency and explainability	Customers and their approvers are not told which parts of the workflow are model-generated, that an invoice without an anomaly flag has not been independently checked, or that their invoices may be used to improve a model serving other ... Recorded by you in the domain assessment.	Write and publish the disclosure and user-notice plan drafted in this assessment, with a named owner and an update process. Document known limitations, including weaker anomaly flagging for suppliers with less than six months of history, and put them where a customer will actually read them.	Dana Whitcombe, Chief Technology Officer	2027-01
High	Human oversight and recourse	Approval is mandatory but is drifting into a formality. Approvers accept prepared payments at a rate that suggests they are not independently checking, and override and rejection rates are not measured, so automation bias would stay ... Recorded by you in the domain assessment.	Report override, rejection and time-to-approve rates per customer each month. Set a floor below which the customer is contacted about approval quality. Give approvers a short briefing on what the model checks and what it does not.	Marcus Oyelaran, Director of Platform Engineering	2027-03

Priority	Domain	Risk and why	Treatment	Owner	Target
High	Legal and regulatory obligations	Customer invoices are used to fine-tune a model that serves other customers, and the customer agreements do not clearly permit it. Sales calls are recorded by a third-party notetaker with no written notice and no data processing agreement. Recorded by you in the domain assessment.	Put the three questions recorded in this assessment to counsel. Build a register of the laws, regulations and contract terms that apply, and keep it current. Stop fine-tuning on customer invoices until the contract position is settled, or obtain explicit written permission from the customers whose invoices are used.	Dana Whitcombe, Chief Technology Officer	2026-12
High	Third-party and vendor dependency	Three AI vendors are in production use with no security assessment: the model provider behind the copilot, the meeting notetaker and the coding assistant. Recorded by you in the domain assessment.	Assess all three AI vendors against a written standard. Request SOC 2 or equivalent from each. Put a data processing agreement in place with the notetaker vendor or remove the tool. Retain current model provider terms and set a review date. Document and test a fallback for the copilot. Add security and AI requirements to the standard supplier contract so this does not recur.	Dana Whitcombe, Chief Technology Officer	2027-02

Priority	Domain	Risk and why	Treatment	Owner	Target
High	Accuracy and reliability	The system misses real cases or raises false alerts at a rate nobody has measured on this deployment. Validation for the intended use is limited or not done. Known limitations you recorded: Extraction degrades on handwritten invoices, photographs of paper, low resolution scans and layouts in languages other than English and French. You noted: Internal target is 98 percent field accuracy on supplier, invoice number, total and bank details, measured against human keying.	Agree acceptance criteria, then measure miss and false-alert rates on your own data, including the conditions listed as known limitations, before relying on the output.	Marcus Oyelaran, Director of Platform Engineering	2027-03
High	Privacy and data protection	Personal information is kept longer than the purpose needs and cannot be reliably deleted. The retention period is unknown.	Set a retention period tied to the purpose, confirm the provider can delete on request, and test deletion once.	Dana Whitcombe, Chief Technology Officer	2027-03
High	Transparency and explainability	The people affected do not know the system is scoring them, what it is used for, or how long records are kept. The disclosure decision is still under review. You noted: There is no plan today, which is why this is under review. What we think it has to cover: a plain-language statement in the product and in the customer agreement that invoice fields are extracted and matched by a model and that anomaly flags are ...	Issue a written notice covering what is collected, how outputs are used, retention, and how to ask questions or complain.	Dana Whitcombe, Chief Technology Officer	2027-01

Priority	Domain	Risk and why	Treatment	Owner	Target
High	Third-party and vendor dependency	The provider reuses or keeps your data for its own purposes, beyond what you intended. Whether the provider may use inputs and outputs is unknown. You noted: Our customer agreements permit processing invoices to deliver the service. They are silent on whether customer invoices may be used to fine-tune a model that serves other customers, which is what we do today.	Amend the contract to prohibit reuse and model training on your data, set a retention limit, and require deletion at exit.	Dana Whitcombe, Chief Technology Officer	2027-02
High	Cybersecurity and misuse	The system or its data is accessed or manipulated through a weakness nobody has tested for. No AI-specific security testing. Access controls are partial or unconfirmed.	Run a security test covering access, data exposure and misuse, and confirm access control and encryption with evidence.	Marcus Oyelaran, Director of Platform Engineering	2027-02
High	Operational resilience and continuity	When the system fails or behaves unexpectedly, nobody knows who acts, how fast, or what the fallback is. No AI incident handling route. Fallback or rollback is partial or not in place.	Name an incident owner, define severity levels and response times, and document the manual fallback when the system is off.	Dana Whitcombe, Chief Technology Officer	Within 30 days
High	Legal and regulatory obligations	The deployment breaches a legal, regulatory or contractual obligation that nobody checked. Legal counsel has not reviewed the use case. You noted: Counsel has not reviewed any of the four AI uses. Three questions need legal input before the ISO 27001 work is signed off.	Have counsel review the use case against privacy, employment and sector obligations and any agreements with affected groups, and record the conditions.	Dana Whitcombe, Chief Technology Officer	2026-12

Priority	Domain	Risk and why	Treatment	Owner	Target
High	Safety and harmful outcomes	A payment released to a fraudulent or altered supplier bank account because extraction read the wrong account number and the anomaly check did not flag it. At our customers typical invoice values that is a direct financial loss, likely unrecoverable, plus a breach of the trust a customer places in us to hold their supplier banking data. Recorded as the most serious foreseeable consequence of failure or misuse.	Define the control that catches this failure before harm, test it, and record who acts when it fires.	Dana Whitcombe, Chief Technology Officer	Within 30 days
Medium	Operational resilience and continuity	Performance drifts or misuse grows without anyone noticing, because nothing sets what is checked, when, or at what threshold. No review frequency is set. Drift monitoring is partial or absent. You noted: CloudWatch alarms exist for latency and error rates on the extraction service and page the on-call engineer.	Set measures, thresholds, a review frequency and an escalation path, and record each review.	Dana Whitcombe, Chief Technology Officer	Within 90 days

Mitigation plan

#	Action	Owner	Target
1	Test error rates across the groups the system affects, record the results, and set a threshold that stops use if the gap is too wide.	Dana Whitcombe, Chief Technology Officer	Before continued use
2	Establish and document the privacy authority for each purpose, including any secondary use, before further processing.	Dana Whitcombe, Chief Technology Officer	2027-03

#	Action	Owner	Target
3	Re-establish a measured accuracy baseline against human keying on a held-out sample, reported monthly per customer and per invoice type. Set an acceptance threshold that blocks a model release. Add alerting on extraction confidence and on the rate of anomaly flags. Produce the approval log extract as standing evidence.	Marcus Oyelaran, Director of Platform Engineering	2027-03
4	Complete a transfer assessment for us-east-1 and decide whether the warehouse moves to ca-central-1 or the transfer is documented and disclosed. Write the internal privacy standard that the public policy implies. Settle in the customer agreement whether invoices may be used to fine-tune a shared model, and stop doing so until that is settled.	Dana Whitcombe, Chief Technology Officer	2027-03
5	Run and record a full privileged access review, then put it on a quarterly cycle. Put leaver revocation on a measurable service level with evidence retained. Turn on alerting for privileged actions and root account use. Bring contractor devices under management or restrict what they can reach. Resume penetration testing and track findings to closure.	Marcus Oyelaran, Director of Platform Engineering	2027-02
6	Write and publish the disclosure and user-notice plan drafted in this assessment, with a named owner and an update process. Document known limitations, including weaker anomaly flagging for suppliers with less than six months of history, and put them where a customer will actually read them.	Dana Whitcombe, Chief Technology Officer	2027-01
7	Report override, rejection and time-to-approve rates per customer each month. Set a floor below which the customer is contacted about approval quality. Give approvers a short briefing on what the model checks and what it does not.	Marcus Oyelaran, Director of Platform Engineering	2027-03
8	Put the three questions recorded in this assessment to counsel. Build a register of the laws, regulations and contract terms that apply, and keep it current. Stop fine-tuning on customer invoices until the contract position is settled, or obtain explicit written permission from the customers whose invoices are used.	Dana Whitcombe, Chief Technology Officer	2026-12
9	Assess all three AI vendors against a written standard. Request SOC 2 or equivalent from each. Put a data processing agreement in place with the notetaker vendor or remove the tool. Retain current model provider terms and set a review date. Document and test a fallback for the copilot. Add security and AI requirements to the standard supplier contract so this does not recur.	Dana Whitcombe, Chief Technology Officer	2027-02
10	Agree acceptance criteria, then measure miss and false-alert rates on your own data, including the conditions listed as known limitations, before relying on the output.	Marcus Oyelaran, Director of Platform Engineering	2027-03
11	Set a retention period tied to the purpose, confirm the provider can delete on request, and test deletion once.	Dana Whitcombe, Chief Technology Officer	2027-03

#	Action	Owner	Target
12	Issue a written notice covering what is collected, how outputs are used, retention, and how to ask questions or complain.	Dana Whitcombe, Chief Technology Officer	2027-01
13	Amend the contract to prohibit reuse and model training on your data, set a retention limit, and require deletion at exit.	Dana Whitcombe, Chief Technology Officer	2027-02
14	Run a security test covering access, data exposure and misuse, and confirm access control and encryption with evidence.	Marcus Oyelaran, Director of Platform Engineering	2027-02
15	Name an incident owner, define severity levels and response times, and document the manual fallback when the system is off.	Dana Whitcombe, Chief Technology Officer	Within 30 days
16	Have counsel review the use case against privacy, employment and sector obligations and any agreements with affected groups, and record the conditions.	Dana Whitcombe, Chief Technology Officer	2026-12
17	Define the control that catches this failure before harm, test it, and record who acts when it fires.	Dana Whitcombe, Chief Technology Officer	Within 30 days
18	Set measures, thresholds, a review frequency and an escalation path, and record each review.	Dana Whitcombe, Chief Technology Officer	Within 90 days

Domain scores

Domain	Weight	Status	Evidence	Residual
Safety and harmful outcomes	12%	Not yet assessed	No evidence provided	100 (default until assessed)
Accuracy and reliability	12%	Assessed	No evidence provided	60
Fairness and human rights	10%	Not yet assessed	No evidence provided	100 (default until assessed)
Privacy and data protection	12%	Assessed	No evidence provided	64
Cybersecurity and misuse	12%	Assessed	No evidence provided	60

Domain	Weight	Status	Evidence	Residual
Transparency and explainability	8%	Assessed	No evidence provided	48
Human oversight and recourse	10%	Assessed	No evidence provided	48
Legal and regulatory obligations	8%	Assessed	No evidence provided	48
Third-party and vendor dependency	8%	Assessed	No evidence provided	64
Operational resilience and continuity	8%	Not yet assessed	No evidence provided	100 (default until assessed)

Monitoring and reassessment

Owner: Marcus Oyelaran, Director of Platform Engineering. Review frequency: Not defined.

Measures and thresholds: Extraction field accuracy sampled monthly against a human-keyed baseline, reported per customer and per invoice type. Duplicate and altered invoices caught before release, compared with the count reported by customers after release. Hours to approval. Copilot deflection rate and the rate of answers a support agent had to correct. None of these are currently measured on a schedule, which is itself a gap we know about.

Second opinion

An independent automated reviewer checked this assessment.

- The vendor domain finding and its risk register entry describe three unassessed AI vendors as the model provider behind the copilot, the meeting notetaker and the coding assistant. The use case scope states explicitly that the coding assistant and the meeting notetaker are internal tools and are not part of the product. The narrative vendor domain section correctly carries all three, but the domain finding scenario does not acknowledge that two of the three sit outside the product boundary. More materially, the compensating control for the vendor domain addresses only the meeting notetaker, while the coding assistant, which has access to the product repository, receives no compensating control. The gap between the stated exposure covering three unassessed vendors and the compensating control covering one vendor partially restricted is not explained.
- The customer equity assessment answer is recorded as not applicable, yet the fairness domain finding, the risk register, and the narrative fairness domain section all identify a genuine fairness exposure: sole proprietor suppliers whose banking details are personal information, and smaller suppliers using handwritten invoices who face higher error rates and payment delays. The narrative correctly notes this tension, stating that the equity assessment was marked not applicable in the intake but the intake itself identifies a fairness exposure. However, the domain

finding does not flag the inconsistency between the customer's own not-applicable answer and the substantive exposure the customer described elsewhere in the same intake. A reviewer reading only the domain finding would not see that the customer's own words contradict that status. This inconsistency could reasonably be surfaced in the domain finding, not only in the narrative.

- The accuracy compensating control states that until a baseline exists, every invoice with changed supplier bank details is queued for two-person approval regardless of anomaly score. The domain finding's own controls section already lists the two-person approval rule for changed bank details as an existing workflow control. The compensating control therefore restates an existing control rather than adding a new or strengthened measure. The human oversight compensating control makes the same point, so the same structural safeguard is presented as a compensating control in two separate domains without explaining what, if anything, is being added beyond what already exists.
- The legal domain compensating control is to suspend the meeting notetaker on customer calls until a data processing agreement and a written notice are in place. The vendor domain compensating control is to disable the meeting notetaker on all customer-facing calls until it has been assessed. Both controls target the same tool and the same population of calls, but they are presented as separate compensating controls under different domains with different expiry dates, December 2026 for legal and February 2027 for vendor. The report does not explain whether these are the same action recorded twice, or two distinct actions with different scope or authority. A reader acting on both would not know which expiry governs.
- The safety domain finding assigns Marcus Oyelaran, Director of Platform Engineering, as owner, and its treatment is to define the control that catches this failure before harm, test it, and record who acts when it fires. The safety risk register entry for the same scenario assigns Dana Whitcombe, Chief Technology Officer, as owner. The two entries describe the same harm pathway but carry different owners with no explanation of the split. A reader following up on the safety domain would contact a different person depending on which section they read.
- The narrative operational resilience domain section states under mitigated that rule-based OCR and outsourced manual keying are noted as available fallbacks for the extraction function if the model is withdrawn. The operational resilience domain finding records no control recorded and the evidence register records no evidence recorded. The narrative draws on the customer alternatives-considered answer, which is a reasonable source, but the domain finding and evidence register do not reflect it at all. A reader of the domain finding alone would not know that any fallback exists, while a reader of the narrative would. The two sections are inconsistent on whether a partial fallback is in place.

Limitations

- 100% of applicable intake questions were completed.
- 0 of 10 domain evidence sets were marked validated and current in the customer workflow.
- A customer-recorded evidence status receives scoring credit only when an independent reviewer and review date are recorded; this is not a RiskWoven authentication of the underlying evidence

unless a separate professional review is completed.

- Narrative answers are not analysed for meaning. Only structural, arithmetic and citation checks are automated; meaning requires qualified human review.
- The advisory output depends on the accuracy and completeness of customer-supplied information.

Advisory disclaimer

RiskWoven AI provides advisory analysis and recommendations only. This report is not legal advice, an audit opinion, a compliance determination, a formal assurance engagement, or a decision on the customer's behalf. The customer remains responsible for decisions, risk ownership, implementation, and qualified professional review.

SPECIMEN