STALWART GOVERNED

AI & Security Governance · vCISO

AI & Security Audit

Riverbend Tax & Books

Engagement ID	Fieldwork	Edition	Status
SG-RVB-2026-A1	2026-08-24 to 2026-08-26	Findings report verifiable source document	FINAL

SAMPLE

This is a sample. Riverbend Tax & Books is a fictional business. Every finding, figure and name in this document was invented to show what a real one looks like. A real client's document is produced by the same method, from evidence gathered at the client.

Pete Saar · AI & Security Governance · vCISO
pete@stalwartgoverned.com · stalwartgoverned.com
Stalwart Governed LLC

© 2026 Stalwart Governed LLC. All rights reserved.

Contents

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

AI & Security Audit

Riverbend Tax & Books

Engagement ID	SG-RVB-2026-A1
Fieldwork	2026-08-24 to 2026-08-26
Edition	Findings report — verifiable source document
Status	FINAL

This is a sample. Riverbend Tax & Books is a fictional business. Every finding, figure and name in this document was invented to show what a real one looks like. A real client's document is produced by the same method, from evidence gathered at the client.

What this is

This report sets out what our AI & Security Audit found, what it means, and what we would do about it, in order of priority. We examined the firm's systems, accounts, vendor terms and day-to-day work directly. A statement made in an interview is never treated as a finding on its own.

Every finding is tied to evidence listed in the Evidence register. Anything we could not confirm is listed under coverage, so the lack of a finding there is never read as a clean result.

This report belongs to the client. It is prepared for their use, and it is theirs to share with their insurers, advisors and counsel, or anyone they choose. We advise having counsel review it before it goes to anyone else outside the firm. The choice is the client's. An unaltered copy of the file can be checked by dropping it on our Verify page, https://stalwartgoverned.com/verify, and we will confirm a copy is genuine on request. This is a security and governance assessment. It is not legal advice. Where it names a statute, a case or a contract term, the client should ask their own counsel how it applies.

© 2026 Stalwart Governed LLC. All rights reserved.

How to read the severity

Five items here need attention today, and each is marked. The rest can be worked through in order. None of this is a judgment about how the firm is run.

Each severity grade is worked out from two inputs: how sensitive the data is (D), and how far outside the firm's control it went (E). We hold ourselves to the same scale. Some findings are about access that stays open, with no data seen leaving through it. These are graded as if data had left. If it is not confirmed where the data went, E is set at E3, the same as a personal or free plan. Each grade is printed with its two inputs beside it. The full scale is printed near the references at the back, so you can see how any grade was figured, and question it.

The bottom line

Five findings: two Critical, three High.

Do today:

- **F-01 — Client tax information pasted into a personal AI chat tool.** Critical (D4×E3) — Tell anyone using the tool to stop pasting client information into it. Turn off any setting that lets it learn from what is typed. This contains the problem until the fix is in place.
- **F-02 — An AI tool that can learn from what is typed, with no data-processing agreement in place.** Critical (D4×E3) — Stop client information going to any AI tool that has no signed data-processing agreement. Turn off learning from what is typed wherever the setting exists. This contains the problem until the fix is in place.
- **F-03 — No second sign-in step on systems that hold client information.** High (D4×E2) — Turn on the platform's own second sign-in step today. Start with the account that holds the most client information. This contains the problem.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

- **F-04 — Access not removed when people leave.** High (D4×E2) — Remove any access still held by someone who has left.
- **F-05 — AI features switched on inside the tax software, with training terms unclear.** High (D4×E2) — List which built-in AI features are switched on in each tool.

For each Critical item: once it is contained, call counsel first, and work through client notice with them. Check your cyber policy for a notice deadline. Telling the insurer is usually a condition of cover. Illinois's breach-notice law, the Personal Information Protection Act (815 ILCS 530), may require notice to affected clients. Counsel decides whether it applies.

Findings

Each piece of evidence below is marked with a code for how it was collected: OBS — seen during a walkthrough · SCR — a screenshot · ADM — a settings export from an admin console · DOC — a vendor document · LOG — a system log entry · FIN — a line on an invoice or statement · INT — an interview, used only as support.

	Finding	Severity
F-01	Client tax information pasted into a personal AI chat tool	**CRITICAL** (D4×E3)
F-02	An AI tool that can learn from what is typed, with no data-processing agreement in place	**CRITICAL** (D4×E3)
F-03	No second sign-in step on systems that hold client information	**HIGH** (D4×E2)
F-04	Access not removed when people leave	**HIGH** (D4×E2)
F-05	AI features switched on inside the tax software, with training terms unclear	**HIGH** (D4×E2)

© 2026 Stalwart Governed LLC. All rights reserved.

F-01 · Client tax information pasted into a personal AI chat tool

CRITICAL (D4×E3) — contain within the hour

What we found.

- `EV-01 · OBS` "a preparer drafts client cover letters by pasting return figures into a free consumer chat tool, though the firm keeps its own shared set of cover-letter templates; the tool's settings let it learn from what is typed" — *walkthrough-notes.md, §1*
- `EV-02 · SCR` "the consumer chat tool's history on the same machine holds three cover-letter drafts with client return figures in them" — *chat-history.png, frame 3*

Why it matters. Pasting return information into a tool the firm does not control may implicate IRC §7216. That section limits when a tax preparer may disclose a client's return information. It is also a gap under the GLBA Safeguards Rule, which requires the firm to oversee the service providers that handle client information. The firm's written information security program does not cover this tool. A personal AI tool may also keep what is typed into it, or learn from it.

What we would do.

- **Now.** Tell anyone using the tool to stop pasting client information into it. Turn off any setting that lets it learn from what is typed. This contains the problem until the fix is in place.
- **Next.** Move the same work into an AI workspace the firm owns, on a business plan whose terms rule out training on what is typed. The firm owns the account, and we only administer it. Add a check-before-you-paste step to the firm's AI use policy (AI & Security Sprint).
- **Keep.** Check it at every monthly review, with one named person responsible for it (AI & Security Retainer).

Referenced: IRC §7216 (criminal) and §6713 (civil); GLBA/FTC Safeguards 16 CFR 314.4(f); the IRS's guidance on a written information security plan (Publications 4557 and 5708); IRS Office of Professional Responsibility, Introductory Guidelines for

© 2026 Stalwart Governed LLC. All rights reserved.

Responsible AI Use in Federal Tax Practice (Issue 2026-19, 24 June 2026); and Circular 230 §10.51(a)(15) · MAP 1.1, MEASURE 2.10 · LLM02

F-02 · An AI tool that can learn from what is typed, with no data-processing agreement in place

CRITICAL (D4×E3) — contain within the hour

What we found.

- `EV-03 · DOC` "the free chat tool's terms, read on the capture date, let it learn from what is typed, and it offers no data-processing agreement at that tier" — *chat-tool-terms.pdf, §4*
- `EV-13 · DOC` "data-processing agreements are on file for both AI services the firm subscribes to, and each says customer data is not used for training" — *dpa-register.pdf, pp. 1–2*

Why it matters. Client information that goes into a vendor's training data cannot be taken back. Without a contract, the firm has no say over what happens to it. That is a gap under the GLBA Safeguards Rule, 16 CFR 314.4(f), which requires the firm to oversee the service providers that handle client information.

What we would do.

- **Now.** Stop client information going to any AI tool that has no signed data-processing agreement. Turn off learning from what is typed wherever the setting exists. This contains the problem until the fix is in place.
- **Next.** Keep client work on AI plans with a signed data-processing agreement that rules out training. Add each AI vendor to the service providers the firm oversees under its written information security program (AI & Security Sprint).
- **Keep.** Re-read each AI vendor's terms at every renewal and whenever the plan changes (AI & Security Retainer).

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Referenced: GLBA/FTC Safeguards 16 CFR 314.4(f) · GOVERN 6.1-6.2 · LLM03, LLM02

F-03 · No second sign-in step on systems that hold client information

HIGH (D4×E2) — same-day

What we found.

- `EV-04 · ADM` "the document-exchange portal's admin console allows sign-in with a password alone; 4 of the 9 users have not set up a second factor" — *portal-security-export.csv, sign-in policy*

Why it matters. Without a second sign-in step, one stolen password is enough to open the client files on those systems. The GLBA Safeguards Rule requires multi-factor sign-in, 16 CFR 314.4(c)(5), so this is a direct gap. In 2022, an insurer sought to void a policy where the application said multi-factor sign-in was in place and it was not.

What we would do.

- **Now.** Turn on the platform's own second sign-in step today. Start with the account that holds the most client information. This contains the problem.
- **Next.** Require a second sign-in step on every system that holds client information (AI & Security Sprint).
- **Keep.** Review who has access, and how they sign in, every quarter (AI & Security Retainer).

Referenced: GLBA/FTC Safeguards 16 CFR 314.4(c)(5); and Travelers v. International Control Services (2022), where the insurer sought to rescind the policy over the application's answer on multi-factor sign-in · MEASURE 2.7

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

F-04 · Access not removed when people leave

HIGH (D4×E2) — same-day

What we found.

- `EV-05 · ADM` "the tax software's user list shows a seasonal preparer whose last day was 2026-04-18 still holds an enabled account" — *tax-users-export.csv, row 7*

Why it matters. A former staff member, contractor or vendor whose access was not removed can still get into client files and tax returns. No one is watching or reviewing that access.

What we would do.

- **Now.** Remove any access still held by someone who has left.
- **Next.** Write a checklist for when someone leaves, and make one named person responsible for it (AI & Security Sprint).
- **Keep.** On every departure, confirm access is removed within a set time and record it (AI & Security Retainer).

Referenced: GLBA/FTC Safeguards 16 CFR 314.4 — the written information security program's required safeguards · MEASURE 2.7

F-05 · AI features switched on inside the tax software, with training terms unclear

HIGH (D4×E2) — same-day

What we found.

- `EV-06 · DOC` "the tax software's built-in AI assistant arrived switched on in a recent update; its terms, read on the capture date, settle that client data is held by the vendor under the firm's existing contract, but do not say whether the assistant learns from it" — *tax-software-terms.pdf, §9*

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Why it matters. Client information is being processed by a third party the firm does not control. The owner cannot yet say how it is used. That is a gap under the GLBA Safeguards Rule, 16 CFR 314.4(f), which requires the firm to oversee the service providers that handle client information.

What we would do.

- **Now.** List which built-in AI features are switched on in each tool.
- **Next.** Review how each tool uses client information, and decide which features stay on. Record each decision with the vendor's data-processing terms (AI & Security Sprint).
- **Keep.** Re-read the terms every year and at every renewal (AI & Security Retainer).

Referenced: GLBA/FTC Safeguards 16 CFR 314.4(f) · GOVERN 6.1 · LLM03

Coverage

This section lists what we examined and what we did not. An area within this audit's scope that we did not examine is listed here, so the lack of a finding there is never read as a clean result.

Area	Status
Backups and tested restore	Examined — in place (EV-07).
Email domain authentication (SPF, DKIM and DMARC)	Examined — in place (EV-10).
Managed security software on computers	Examined — in place (EV-08).
AI browser extensions and client information	Examined — in place (EV-14).

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Area	Status
Qualified Individual and written information security program	Examined — in place (EV-11).
Software updates	Examined — in place (EV-09).
Periodic review of third-party access	Examined — in place (EV-12).
AI note-takers on client calls	Described in an interview, and not confirmed. It is recorded as a lead only. To check it: list the apps connected to the calendar and the meeting software, and see which of them join client calls.

A row marked "in place" was examined, and the control was found working. The row names its evidence, so you can check it in the Evidence register. Every other row is a place to look next. None of those rows is a finding, and none is graded, because nothing there was confirmed.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Common cyber-insurance questions

The rows below are questions cyber-insurance applications commonly ask, with what this audit found for each during fieldwork. They are not any one insurer's form, and this table does not answer an application for the firm. If the firm gives us its current application, this section uses that form's questions.

Question	What this audit found during fieldwork
A second sign-in step (multi-factor authentication) on email, remote access and admin accounts	Not in place where examined (F-03). Remote access and admin accounts are not part of this review. To check the rest: export the sign-in settings for email, remote access and each admin account, and see which accept a password alone.
Security software on every computer that detects and responds to attacks (EDR)	Examined in part. Managed security software on computers: in place (EV-08). This review does not test whether the software detects and responds to attacks, beyond protecting each computer. To check the rest: see in the security software's console whether its detection and response features are switched on, and who is alerted when it finds something.
Backups kept offline or apart from the network, with restores tested	Examined in part. Backups and tested restore: in place (EV-07). This review does not test whether a copy is kept offline or apart from the network. To check the rest: see where each backup copy is stored, and whether it can be changed or deleted from the office network.
Software updates applied, and no systems past their end of support	Examined in part. Software updates: in place (EV-09). End-of-support systems are not part of this review. To check the rest: list any computer or program whose maker no longer issues updates.
Training staff to spot phishing emails	Not examined. To check it: read the training records and the results of any phishing test the office has run.
Confirming any request to move money or change payment details by a second channel (funds-transfer verification)	Not examined. To check it: walk through how the office confirms a request to send money or change payment details.

© 2026 Stalwart Governed LLC. All rights reserved.

Question	What this audit found during fieldwork
A written plan for responding to an incident (incident response plan)	Not examined. To check it: see whether there is a written incident response plan, and whether it covers AI tools and when a breach must be reported.

A row marked "Not examined" means this audit holds no evidence either way. It is not a clean result.

© 2026 Stalwart Governed LLC. All rights reserved.

Evidence register

Every piece of evidence in this engagement, in order, and where it is used. Nothing collected is left out. An item that could not be graded is listed under Coverage. Each item's fingerprint is the first 16 characters of its source file's SHA-256, taken when the file was collected. To check a copy of a collected file, compute its SHA-256 and compare the first 16 characters. The full SHA-256 is kept with the engagement record.

	Code	Source	Fingerprint (first 16)	Used in
EV-01	OBS	walkthrough-notes.md, §1	`cb40d533431235e6`	F-01; V-03
EV-02	SCR	chat-history.png, frame 3	`a2a61cdcc0658f10`	F-01
EV-03	DOC	chat-tool-terms.pdf, §4	`dd78ef3a631ce37d`	F-02
EV-04	ADM	portal-security-export.csv, sign-in policy	`057ad2aaa946eb16`	F-03
EV-05	ADM	tax-users-export.csv, row 7	`e8e1c5233fade2bb`	F-04
EV-06	DOC	tax-software-terms.pdf, §9	`f981df70b728d0e3`	F-05
EV-07	LOG	backup-restore-log.txt, 2026-07-28	`53e31df5bede3afe`	Coverage — Backups and tested restore (in place)
EV-08	ADM	endpoint-console.png, fleet view	`59fb70c7d3ff4381`	Coverage — Managed security software on computers (in place)
EV-09	ADM	update-console.csv, all rows	`301af9b6cf0c746f`	Coverage — Software updates (in place)

© 2026 Stalwart Governed LLC. All rights reserved.

	Code	Source	Fingerprint (first 16)	Used in
EV-10	ADM	dns-records.txt, SPF · DKIM · DMARC	`fadbcf07fd2edda9`	Coverage — Email domain authentication (SPF, DKIM and DMARC) (in place)
EV-11	DOC	wisp-v3.pdf, §1	`e0e35c063276e57d`	Coverage — Qualified Individual and written information security program (in place)
EV-12	ADM	access-review-q2.pdf, p. 1	`b0c5ccf458bdd888`	Coverage — Periodic review of third-party access (in place)
EV-13	DOC	dpa-register.pdf, pp. 1–2	`6f0b24e8987a6ec2`	F-02
EV-14	ADM	extension-inventory.csv, all rows	`a8c388eef108f4e1`	Coverage — AI browser extensions and client information (in place)
EV-15	INT	interview-notes.md, §4	`7c8ba50d7ec924c8`	Coverage — AI note-takers on client calls
EV-16	FIN	bookkeeping-invoice-aug.pdf, line 3	`894d113cc284af08`	V-01
EV-17	ADM	assistant-usage.csv, 30-day view	`5eade07304d293d9`	V-01
EV-18	ADM	assistant-billing.png, cycle view	`d8e43ddefe75c194`	V-02
EV-19	DOC	ai-use-policy.pdf, p. 1	`667eacf0e2d923b0`	V-04
EV-20	SCR	prompt-library.png, version history	`30ebde5e1a263105`	V-05
EV-21	OBS	walkthrough-notes.md, §6	`cb40d533431235e6`	V-06

© 2026 Stalwart Governed LLC. All rights reserved.

How the firm is using AI

The findings above cover risk. This section covers value: what the firm pays for and what it uses, where its tools are set up well, and where effort is wasted. It also shows where one practice brings both value and risk. Every row traces to the same evidence as the findings. Where the evidence did not include a figure, it is listed under "What the evidence did not include". No figure is estimated.

What the firm pays for and what it uses

This part weighs what the firm pays for against what it uses. It covers seats paid for against seats in use, credits bought against credits used, firm money spent on personal AI plans, and jobs paid for twice because two tools do them. The figures come from the vendors' and the firm's own pages, named below. Nothing here is calculated from them.

	Reading	Evidence
V-01 · Paid AI seats against active users	Under-used	EV-16 · EV-17
V-02 · Credits used against the plan's allowance	Headroom	EV-18
V-03 · Two or more tools doing the same job	Costing effort	EV-01

V-01 · Paid AI seats against active users

Under-used — 4 seats paid for, 3 active in the period.

What we found.

- `EV-16 · FIN` "the bookkeeping suite's invoice carries its assistant add-on on 4 seats, every bookkeeping seat the firm holds" — *bookkeeping-invoice-aug.pdf, line 3*
- `EV-17 · ADM` "the assistant add-on's usage report shows 3 of the 4 seats active in the trailing 30 days" — *assistant-usage.csv, 30-day view*

© 2026 Stalwart Governed LLC. All rights reserved.

What this shows. A seat that is paid for and not signed into is a recurring charge for nothing. A seat used by more than one person is a shared login, which is a separate finding. The license page and the usage report, read together, show what the firm gets for that line of the invoice.

V-02 · Credits used against the plan's allowance

Headroom — 300 credits in the cycle's allowance, 210 used.

What we found.

- `EV-18 · ADM` "the assistant add-on's billing dashboard shows an allowance of 300 assistant credits in the cycle, with 210 used at capture" — *assistant-billing.png, cycle view*

What this shows. The plan's allowance is what the firm bought. The credits used are what it used. Credits left over every cycle point to a smaller plan. A cycle that runs over points to a bigger plan, or to a workflow worth reviewing. The figures are the vendor's own, read off its dashboard.

V-03 · Two or more tools doing the same job

Costing effort — effort is being wasted here. It is recorded as what was seen, without a grade.

What we found.

- `EV-01 · OBS` "a preparer drafts client cover letters by pasting return figures into a free consumer chat tool, though the firm keeps its own shared set of cover-letter templates; the tool's settings let it learn from what is typed" — *walkthrough-notes.md, §1*

What this shows. One job is done by more than one tool. Each tool has its own account, its own terms and its own place the work is sent. Choosing one tool and moving

© 2026 Stalwart Governed LLC. All rights reserved.

the work onto it is usually the cheapest consolidation a firm this size can make. It also shortens the lists in this report.

What works well and where effort is wasted

What the firm does well is recorded here the same way as a finding: seen, cited and named. Beside it are the places where the time a tool saves is lost again.

	Reading	Evidence
V-04 · A written AI use policy, acknowledged by staff	Written, not followed	EV-19
V-05 · A shared library of prompts and templates, with a record of changes	Working well	EV-20
V-06 · AI built into a workflow, with a person reviewing the output first	Working well	EV-21

V-04 · A written AI use policy, acknowledged by staff

Written, not followed — the policy exists on paper, and the work seen during the same fieldwork does not follow it. It is recorded so a written policy is never mistaken for protection.

What we found.

- `EV-19 · DOC` "a one-page AI use policy dated 2026-02-02 is acknowledged in writing by all nine staff, but this audit's walkthrough found a preparer pasting client return figures into a free chat tool, which the policy forbids" — *ai-use-policy.pdf, p. 1*

What this shows. A signed policy that staff do not follow does not work as a control. An insurer may read it as a promise the firm did not keep. The usual fix is to name the approved tools, make the approved way the easy one, and check that it is the one being used.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

V-05 · A shared library of prompts and templates, with a record of changes

Working well — seen operating as intended during the same fieldwork. It is recorded as a strength and cited the same way as a finding. It does not provide assurance.

What we found.

- `EV-20 · SCR` "a shared folder holds a versioned set of client-letter prompt templates, with version history on and edits by three staff in the past quarter" — *prompt-library.png, version history*

What this shows. When more than one person edits a prompt library, skill becomes something the whole firm can use. It no longer sits in one person's chat history. It is the plainest sign that the firm keeps what it learns about the tool.

V-06 · AI built into a workflow, with a person reviewing the output first

Working well — seen operating as intended during the same fieldwork. It is recorded as a strength and cited the same way as a finding. It does not provide assurance.

What we found.

- `EV-21 · OBS` "the bookkeeper runs month-end reconciliation exceptions through the assistant and reads each suggested match before posting it; nothing posts unread" — *walkthrough-notes.md, §6*

What this shows. The tool sits inside a real workflow, and a person reads its output before it takes effect. Every good AI workflow works this way. We point it out when we find it, so the rest of the firm can copy it.

What the evidence did not include

Everything this part of the assessment covers had evidence behind it. Nothing is missing.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Value against risk

Some rows above cite the same evidence as a finding. In those cases, one practice both saves time and carries a risk. The fix for that finding can keep the time saved. Where a row has no finding beside it, no finding cites its evidence.

	Reading	Finding on the same evidence
V-01 · Paid AI seats against active users	Under-used	—
V-02 · Credits used against the plan's allowance	Headroom	—
V-03 · Two or more tools doing the same job	Costing effort	F-01 Critical (D4×E3)
V-04 · A written AI use policy, acknowledged by staff	Written, not followed	—
V-05 · A shared library of prompts and templates, with a record of changes	Working well	—
V-06 · AI built into a workflow, with a person reviewing the output first	Working well	—

Nothing in this section is a grade. It records what the firm pays for, what it uses and what is worth keeping, so the firm can make those choices on purpose.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

The severity scale

The scale is printed here so every grade can be checked. D is how sensitive the data is. E is how far outside the firm's control it went. Each finding states its own two inputs below. Some findings are about access that stays open, with no data seen leaving through it. These are graded as if data had left. If it is not confirmed where the data went, E is set at E3, the same as a personal or free plan. For a control that is missing or switched off, E is graded by what the control protects. A sign-in control that is switched off is graded by the system it guards, whether or not any data was seen leaving.

	E1 — stayed on the firm's own machines or in an AI workspace the firm owns	E2 — a business plan the firm controls	E3 — a personal or free plan, or an unconfirmed destination	E4 — confirmed outside the firm, or the tool took an action on its own
D4 — regulated data, sign-in credentials, or access keys and tokens	MEDIUM	HIGH	CRITICAL	CRITICAL
D3 — confidential client business information	LOW	MEDIUM	HIGH	CRITICAL
D2 — internal or anonymized information	LOW	LOW	MEDIUM	HIGH
D1 — public information	LOW	LOW	LOW	MEDIUM

- **F-01** — D4 (regulated data, sign-in credentials, or access keys and tokens) × E3 (a personal or free plan) → Critical

© 2026 Stalwart Governed LLC. All rights reserved.

- **F-02** — D4 (regulated data, sign-in credentials, or access keys and tokens) × E3 (a personal or free plan) → Critical
- **F-03** — D4 (regulated data, sign-in credentials, or access keys and tokens) × E2 (a business plan the firm controls) → High
- **F-04** — D4 (regulated data, sign-in credentials, or access keys and tokens) × E2 (a business plan the firm controls) → High
- **F-05** — D4 (regulated data, sign-in credentials, or access keys and tokens) × E2 (a business plan the firm controls) → High

---

References

Every framework code cited in the findings above, and the document it names. Statutes and regulations are named in full where they are cited, in each finding's own Referenced line.

- **MAP 1.1** — NIST AI Risk Management Framework 1.0: "Intended purposes, potentially beneficial uses, context-specific laws, norms and expectations, and prospective settings in which the AI system will be deployed are understood and documented." airc.nist.gov/AI_RMF_Knowledge_Base/AI_RMF/Core_And_Profiles/5-sec-core
- **MEASURE 2.10** — NIST AI Risk Management Framework 1.0: "Privacy risk of the AI system – as identified in the map function – is examined and documented." airc.nist.gov/AI_RMF_Knowledge_Base/AI_RMF/Core_And_Profiles/5-sec-core
- **LLM02** — OWASP Top 10 for Large Language Model Applications, LLM02:2025 Sensitive Information Disclosure: "Sensitive information can affect both the LLM and its application context. This includes personal identifiable information (PII), financial details, health records, confidential business data, security credentials, and legal documents." genai.owasp.org/llmrisk/llm022025-sensitive-information-disclosure/

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

- **GOVERN 6.1-6.2** — NIST AI Risk Management Framework 1.0: "GOVERN 6.1 — Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights. GOVERN 6.2 — Contingency processes are in place to handle failures or incidents in third-party data or AI systems deemed to be high-risk." nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
- **LLM03** — OWASP Top 10 for Large Language Model Applications, LLM03:2025 Supply Chain: "LLM supply chains are susceptible to various vulnerabilities, which can affect the integrity of training data, models, and deployment platforms. These risks can result in biased outputs, security breaches, or system failures." genai.owasp.org/llmrisk/llm032025-supply-chain/
- **MEASURE 2.7** — NIST AI Risk Management Framework 1.0: "AI system security and resilience – as identified in the map function – are evaluated and documented." airc.nist.gov/AI_RMF_Knowledge_Base/AI_RMF/Core_And_Profiles/5-sec-core
- **GOVERN 6.1** — NIST AI Risk Management Framework 1.0: "Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights." nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Sign-off

Prepared by Peter Saar, CISSP #2450556.

AI & Security Governance · vCISO.

Issued by Stalwart Governed LLC.

This report records what was observed during the fieldwork window named above. It is not a certification, and nobody can certify that a business is safe. It is evidence, and a set of recommendations that follow from it.

STALWART GOVERNED
AI & Security Governance · vCISO

Signed: Peter Saar · Date of issue: 2026-08-28

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.

Record

This is a sample, and you can check it the way a client would: drop this file on our Verify page, https://stalwartgoverned.com/verify, and it will show as a sample we published, unchanged. It will never show as an issued document. The fingerprint in the first row is of this sample's source document, the file the row names. It is not the fingerprint of this PDF. The PDF has its own, and both are registered. On a real engagement, the source document is delivered with the PDF, so either file can be checked on our Verify page. The remaining rows identify this document's version and date the reference pack it cited.

sha256 of riverbend.md	7c85a72dd35fd03be367f03d32ec4e360ffb0877d6ea2221e78edaef2367f2eb
length	365 lines · 30162 bytes
copy	1 of 3
version	d40fd37c5203 · cd154950a3cb
reference pack	2026-08-21
to check it	certutil -hashfile "riverbend.md" SHA256 (Windows) · shasum -a 256 "riverbend.md" (macOS) · sha256sum "riverbend.md" (Linux)

How this report was made

The findings, grades and recommendations in this report were decided by the firm. Each finding rests on evidence the firm collected, each grade is computed by the firm's fixed method, and each figure is read directly from the evidence. We used AI as a drafting tool within fixed rules for what a sentence may and may not say. In signing the report, we take full editorial responsibility for it and adopt the words as our own.

SAMPLE · Riverbend Tax & Books is a fictional business
© 2026 Stalwart Governed LLC. All rights reserved.