

**AIGAS**

AI Governance Assurance Standard

Top 5 AI risk areas in accounting practice

A practical guide to protecting client data and professional judgement

AIGAS · SEPTEMBER 2026 · VERSION 1.0

AI can help a practice draft correspondence, organise information and reduce repetitive work. The difficulty is knowing whether those time savings come with errors, disclosure of client information or actions that nobody intended to authorise.

This guide helps partners, practice managers and team leaders examine five areas where everyday AI use can weaken familiar controls. It provides practical steps, worked scenarios and reusable checks for tax, accounts, bookkeeping and advisory teams, with additional points for audit firms.

Start with these five questions

Risk area	Question for the firm
1 Shadow AI	Can we identify the tools, accounts and AI features staff actually use?
2 Client data exposure	Do we know what information goes in and who can access it?
3 Unreliable outputs	Can we support the answer with appropriate sources and calculations?
4 Ineffective human review	Can the reviewer detect an error and stop the work being released?
5 Uncontrolled automation	Can AI change, send or submit anything beyond its approved authority?

Use this guide in a partner meeting

Choose one real workflow, such as drafting tax advice or reviewing management accounts. Trace it through all five questions. Assign an owner to each gap, then use the 30-day plan on page 9. Start with live client-data exposures and uncontrolled actions before spending time polishing policy wording.

These are AIGAS priority areas, not a statistical ranking or an exhaustive risk assessment. The scenarios are illustrative, not reported incidents. The controls and templates are AIGAS recommendations informed by the sources on page 10; they are not an ICAEW or AccountingWEB checklist. This guide does not replace applicable professional standards or engagement-specific judgement.

1 Shadow AI

The risk is work taking place through tools, accounts or features that the firm has not assessed or cannot see.

The analogy

Imagine a member of staff quietly hiring an assistant to help with client work. The assistant may be useful, but nobody has checked their contract, access or competence. Unapproved AI creates a similar gap in the firm's oversight.

How it appears in practice

A team member uses a personal chatbot account to summarise correspondence. Another installs a browser writing extension. A meeting assistant joins a client call, while an AI feature becomes available inside existing practice software. None appears on the firm's approved software list because no new supplier invoice has reached the partner.

ICAEW's May 2026 audit guidance expressly identifies the risk of staff transferring personal AI habits into professional work without firm authorisation or oversight. Its scope is audit; the discovery problem is relevant to other practice workflows too. [1]

Warning signs

Staff say "we only use Copilot" but cannot specify the account or configuration. Personal email addresses appear in work-related AI accounts. The policy lists a few chatbot brands but says nothing about transcription, extensions or embedded features.

Actions to take this week

- Ask every team, including partners and contractors, what AI they use, for which tasks, through which account, and with what data. Explain that the purpose is to establish safe working arrangements so people can answer candidly.
- Check the responses against business subscriptions, installed extensions, meeting tools and connected applications with your IT provider. Use proportionate, transparent checks; staff declarations alone will miss some usage.
- Record the product, exact plan or tenant, feature, task, owner, users, data permitted and approval status. Approval for generic drafting should not silently become permission to upload payroll records.
- Give staff an accessible approved route and a named person for new requests. Pending assessment, restrict unknown tools to public or invented material. Pause any use involving live client data where protection cannot be established.
- Review the inventory when people join or leave, when a new feature is enabled and when a supplier changes relevant terms. Put a regular review date in the diary.

Evidence and a practical test

Keep the dated inventory, approval decisions and unresolved requests. Ask a team member to demonstrate a recent workflow using non-client material. Does the actual account and feature match the register? A brand name alone is insufficient evidence.

Owner: practice manager or nominated AI lead, with partner accountability. First measure: unexplained tools or accounts still in use, and how long approval requests remain unresolved.

2 Client data exposure

The risk is client information being disclosed, retained or reused beyond the firm's approved arrangements, including inside its own systems.

The analogy

Uploading a client file to an unassessed AI service is like sending it to a new subcontractor before checking who they are allowed to share it with. Paying for the service does not answer that question.

How it appears in practice

A manager uploads a payroll spreadsheet to explain an unexpected movement. It includes names, salaries and bank details although only departmental totals were needed. Elsewhere, an internal assistant can retrieve a sensitive transaction file because the underlying folder permissions are too broad.

ICAEW's July 2026 confidentiality reminder covers internal access as well as external disclosure. It calls for appropriate due diligence, contractual protections and internal approval before confidential information enters digital or AI systems. [2]

Warning signs

The firm relies on "enterprise", "secure" or "not used for training" as its whole assessment. Staff remove a client name but leave identifying transactions. Nobody can explain retention, support access, subprocessors or connected-file permissions.

Actions to take this week

- Classify permitted data for each use case. Separate public material, confidential business information, personal data and especially sensitive records. Use the smallest extract needed; test with invented data first.
- Check the exact service and plan. Establish contractual data use, model-training arrangements, storage and processing locations, retention and deletion, support access, subprocessors, security and incident notification. Record gaps and restrict use until resolved.
- Have the appropriate lead assess confidentiality authority, engagement terms and applicable data-protection duties. These are separate questions: client permission alone does not settle every legal requirement. Obtain specialist advice where needed.
- Restrict file and connector access to the people and tasks that need it. Test whether a user can retrieve a document from another engagement. Review shared chat links, exports and meeting recordings as well as prompts.
- Teach staff that removing names may leave a client identifiable from context. Confirm that redaction removes hidden spreadsheet sheets, comments and document metadata where relevant; do not assume a shortened file is safe.

Evidence and a practical test

Keep the supplier assessment, approved data categories, configuration record and access-test results. Pick one workflow and trace the information from the original file through prompts, outputs, storage and deletion. Who can see each copy?

Owner: the firm's data-protection or confidentiality lead, supported by IT and the engagement partner. If disclosure is suspected, stop further uploads and use the incident steps on page 9. Deleting a chat does not by itself establish that every retained copy has gone.

3 Unreliable outputs

The risk is an incorrect, incomplete or unsuitable answer entering professional work because it appears coherent and credible.

The analogy

Treat an AI answer like a polished set of accounts with no working papers. Good presentation makes it readable; it does not substantiate the figures or conclusions.

How it appears in practice

An AI draft uses the wrong tax year, overlooks a client-specific exception or cites a source that does not support the conclusion. An accounts summary reports a movement correctly but invents its cause. A spreadsheet formula gives the expected answer on one row and fails on credits or blank cells.

The problem also arrives from clients. AccountingWEB reported practitioners spending time correcting AI-generated tax claims and dealing with the resulting pressure on fixed fees. [3] Separately, its coverage of financial statement review raises the limitations of using general-purpose AI for precision work. [4]

Warning signs

A confident answer has no inspectable sources. A citation exists but concerns another jurisdiction. Totals reconcile while underlying classifications are wrong. A reviewer asks a second chatbot to confirm the answer and treats agreement as verification.

Actions to take this week

- State the task's boundaries before use: jurisdiction, period, entity, facts, reporting framework and missing information. Ask for assumptions and uncertainties to be identified, then check them yourself.
- Open material sources independently. Confirm the cited document exists, is current for the relevant period and supports the claim. Use authoritative material and appropriate technical research resources; AI-generated citations are leads to investigate.
- Repperform material calculations with an appropriate calculation tool and reconcile inputs to source records. For generated spreadsheet logic, test normal cases, zeros, negatives, blanks and exceptions before wider use.
- Check completeness as well as correctness. Compare the output with the engagement's requirements and relevant disclosure or technical checklists. A correct answer to the wrong or incomplete question can still mislead.
- For client-supplied AI claims, establish the decision the client is trying to make and the actual facts first. Ask for the original answer and sources where useful. Agree additional scope before extensive investigation outside the engagement.

Evidence and a practical test

Retain relevant sources, calculation checks, material corrections and the reasoned conclusion within the controlled engagement file. A prompt history alone does not demonstrate correctness. Avoid duplicating unnecessary confidential material in a separate AI log.

Owner: the technical lead or engagement reviewer. Test one recently completed AI-assisted job: can another competent accountant reconstruct the basis for its key conclusion without relying on the chatbot?

4 Ineffective human review

The risk is apparent oversight that cannot identify or challenge a problem. Risk 3 concerns the answer; this area concerns whether the review process can catch its defects.

The analogy

A second signature on a reconciliation means little if the signatory has not seen the bank statement. “Human reviewed” is equally weak when the person lacks the source records, expertise or time to check the result.

How it appears in practice

A junior prepares an AI-assisted advice note. A busy manager reads it for tone and assumes technical checking has already happened. The junior assumes the manager will verify the tax treatment. The firm has two people in the workflow, but neither has performed the critical check.

PCRT guidance on ethical AI use emphasises competence, professional scepticism and continuing accountability for work involving AI. It also addresses AI-generated client information and review of automated tax-filing workflows. [5]

Warning signs

Every workflow uses the same “reviewed” tick box. Training covers prompt writing but not errors. Reviewers cannot explain the basis of an answer in their own words. Time savings are measured before rework and review time are counted.

Actions to take this week

- Define review requirements by consequence. Routine internal drafting may need a proportionate check; tax positions, financial statements and client-facing advice require a reviewer competent in the subject and the engagement.
- Specify what is checked: facts, sources, calculations, completeness, professional judgement and intended recipient. Make clear who prepares, who verifies and who authorises release. In a small firm, record how these functions are performed when roles overlap.
- Give reviewers the input records, material assumptions and tool limitations. Allocate enough time and explicit authority to reject output or revert to the established process.
- Use a short exercise with an invented client and deliberately flawed output. Include a wrong period, missing exception and plausible but unsupported statement. Debrief what was missed; adapt training to demonstrated needs.
- Keep technical learning in the workflow. Ask junior staff to explain conclusions and show a manual check. Monitor whether AI is removing opportunities to develop the knowledge needed to review future work.

Evidence and a practical test

Record who reviewed the work, which checks they performed, exceptions raised and how those exceptions were resolved. Keep training records alongside evidence that staff can apply the controls. Completing a course is not the same as demonstrating competence.

Owner: the engagement partner and training lead. Ask the named reviewer to explain one significant conclusion without reading the AI text. If they cannot, reassign the review or narrow the use case. For sole practitioners, identify when an external technical consultation or independent check is needed.

5 Uncontrolled automation

The risk is AI taking actions or gaining access beyond the authority intended by the firm, or a workflow changing after it was approved.

The analogy

Giving a new assistant the keys to the filing cabinet is one decision. Giving them authority to change bank details and send client correspondence is another. AI access and AI authority need the same distinction.

How it appears in practice

A tool approved to draft replies is connected to a mailbox with permission to send. An invoice workflow can update accounting records before an exception is reviewed. A supplier update introduces a new action or changes how information is processed, but the original approval remains untouched.

External content can also influence an AI system. The NCSC describes indirect prompt injection: material encountered by an AI system is treated as an instruction. [6] For a practice, the consequence depends on what the connected system can access or do.

Warning signs

An integration requests broad permissions with no documented reason. Nobody owns supplier changes. There is no tested way to disable the workflow or establish which records it changed. A successful demonstration is treated as sufficient testing.

Actions to take this week

- List each permission separately: read, create a draft, edit a record, send, submit or trigger a payment. Begin with the minimum access and authority needed. A supplier's capability should not automatically become the firm's permission.
- Retain explicit human approval before consequential actions, such as changing payment details, issuing advice or submitting a return, unless a separately assessed and tested control design justifies otherwise. Keep existing independent payment-verification procedures.
- Test with invented records and adverse cases: a misleading attachment, wrong recipient, duplicate invoice, unavailable service and unexpected instruction inside a document. Establish what should stop the workflow and who receives the exception.
- Name a supplier and change owner. Reassess significant changes to models, features, permissions, terms and workflows. Use a small, known set of example tasks to check behaviour after material changes.
- Document how to suspend the workflow, revoke access, inspect action logs and restore or correct affected records. Test a manual fallback so a service outage does not turn into missed client deadlines.

Evidence and a practical test

Keep the permission record, acceptance tests, change decisions and action logs appropriate to the risk. Record who can operate the stop procedure and where the instructions are held. Logs should be access-controlled and retained proportionately.

Owner: the workflow owner with IT support and partner approval. In a controlled test, ask them to stop the automation and identify everything it changed. If they cannot, restrict it to a safer level of access until the gap is resolved.

A quick review for your next partner meeting

Work through these ten checks against actual evidence. Record Yes, Partly, No or Unknown. Do not turn the answers into a certification score: one serious exposure matters more than several completed policies.

Area	Checks to answer	Status or action
Shadow AI	1. Our inventory matches sampled real workflows. 2. Each live use has an owner and defined approval boundaries.	
Client data	3. We know the data terms and settings of approved services. 4. We have tested access restrictions and data minimisation.	
Outputs	5. Material claims and calculations receive independent checks. 6. We check completeness and client-specific assumptions.	
Human review	7. Reviewers have the competence, information and time required. 8. Work records show the checks and resolution of exceptions.	
Automation	9. Permissions and consequential actions are explicitly controlled. 10. We can stop the workflow and identify affected work.	

How to prioritise the findings

Act immediately where live client data enters an unassessed service, a material conclusion cannot be supported, or an automation can take an uncontrolled consequential action. Pause the affected activity while a responsible person assesses it.

Next, resolve “Unknown” answers in active workflows. Unknown does not prove a breach or an error; it means the firm lacks the information needed to rely on its control. Assign a named owner and a date for establishing the facts.

Then address gaps in evidence, training and routine monitoring. A policy in a shared folder is like a fire procedure nobody has practised: it can describe the intended response without establishing that people can perform it.

Three useful management measures

- Visibility: live uses with unresolved approval or data questions, tracked to closure.
- Quality: material errors found during review and after release, with recurring causes.
- Value: net time saved after verification and rework, alongside incidents and near misses.

Review these measures together. Faster production is valuable only if the practice can still substantiate the work and protect its clients.

Reusable records for everyday work

A use case approval record

Copy these fields into your existing governance system or register. Keep supporting contracts and engagement records in their appropriate controlled locations.

Field	What to record
Use and ownership	Task; team; accountable owner; product; plan or tenant; AI feature; date.
Boundaries	Permitted inputs; prohibited data; users; outputs; actions allowed; actions requiring approval.
Supplier and access	Assessment location; relevant terms and settings; connectors; permissions; unresolved questions.
Quality and review	Test cases; acceptance criteria; named reviewer role; verification required; escalation route.
Decision and monitoring	Approved, restricted or paused; decision maker; conditions; review date; change triggers; stop procedure.

Worked example of a restricted approval

Task: draft an explanation of movements in management accounts. Inputs: invented test data initially; an approved minimum client extract only after the data assessment. Output: internal draft commentary. Authority: no ledger edits and no client sending. Reviewer: manager checks figures to the approved accounts and confirms each proposed explanation against client evidence. Unexplained movements remain open questions. Retest after a material supplier or workflow change.

An engagement review note

Use this short structure for work where AI materially contributes to a conclusion or deliverable. Adapt the detail to the engagement and applicable record-retention requirements.

- Work reference and AI use: task, tool, date and material inputs or source-record references.
- Verification: sources consulted, calculations reperformed and completeness checks.
- Judgement: assumptions challenged, missing facts requested and reasoning supporting the conclusion.
- Exceptions: errors identified, corrections made, unresolved issues and escalation.
- Release: reviewer, date and authorisation to use the result.

A response when a client brings an AI answer

“Thank you for sending this. Before we rely on it, we need to check the facts, relevant period and sources against your circumstances. Please send the original answer and confirm the decision you are considering. We will explain the applicable position and agree any additional work outside our current engagement before proceeding.”

A practical 30 day action plan

The timetable below is a starting point for a small or medium-sized practice. Combine roles where necessary, but retain a named accountable partner. Deal with an active exposure immediately rather than waiting for its scheduled week.

When	Action and owner	Evidence of completion
Days 1 to 5	Partner appoints a lead. Practice manager discovers current uses. Pause unassessed client-data use and uncontrolled consequential actions.	Initial inventory, immediate restrictions and named owners.
Days 6 to 10	Data lead and IT assess priority suppliers, settings and permissions. Team heads select approved use cases.	Use-specific approvals, data boundaries and outstanding issues.
Days 11 to 20	Technical and training leads define review requirements and run a flawed-output exercise. Workflow owners test automation limits.	Completed review examples, training actions and test results.
Days 21 to 30	Partner samples live work, checks unresolved risks and compares gross savings with review and rework.	Recorded decisions, action dates and next review meeting.

If something has already gone wrong

1. Contain the affected activity. Stop uploads, restrict access or suspend the automation as appropriate. Preserve relevant evidence in a controlled location; avoid deleting records needed to establish what happened.
2. Establish the scope. Record the tool and account, data involved, recipients, actions, times, affected engagements and decisions that relied on the output. Involve the relevant partner, IT and data-protection lead promptly.
3. Assess required notifications and corrective action. The responsible lead should determine applicable client, contractual, insurer, professional and regulatory obligations and deadlines, obtaining specialist advice where needed. Do not wait for a routine review meeting.
4. Correct the work and verify recovery. Identify other jobs affected by the same fault, recheck outputs and document the decision to resume. Amend the control that failed and share the lesson without unnecessarily exposing client information.

Start by making AI use visible

AIGAS provides a practical AI governance framework for accounting firms. The free AI Tool Register offers a starting point for recording AI use across the practice. An inventory helps you begin; it must be followed by decisions, appropriate controls and evidence that those controls work.

[Start the free AI Tool Register](#)

[Explore AIGAS](#)

Sources and further reading

Research reviewed on 29 September 2026. Professional guidance and supplier arrangements can change. Read the current source when making an engagement-specific decision. References support the issues identified; AIGAS's examples, suggested controls and templates are its own practical synthesis.

[1] ICAEW | Artificial intelligence in audit work managing the risks | 6 May 2026

Professional Standards Department guidance on unauthorised use, training, oversight and audit quality. Audit-specific expectations should not be presented as rules for every service.

[Read source 1](#)

[2] ICAEW | Protecting client confidentiality internal controls and the ICAEW Code of Ethics | 28 July 2026

Explains confidentiality within and outside the firm and considerations before using confidential information in AI tools.

[Read source 2](#)

[3] AccountingWEB | Clients' use of AI is wasting accountants' time | 10 September 2025

Richard Hattersley reports practitioners' experience of client-generated misinformation and the associated time and fee pressures.

[Read source 3](#)

[4] AccountingWEB | Reviewing financials Why ChatGPT isn't enough | 2026

Coverage of precision in financial statement review, including a software supplier perspective. This is contextual reporting, not a comparative product evaluation.

[Read source 4](#)

[5] ICAEW and PCRT bodies | Topical guidance covering the application of PCRT to the ethical use of artificial intelligence tools | Current online guidance

Tax-focused ethical guidance on competence, accountability, objectivity, confidentiality and review.

[Read source 5](#)

[6] National Cyber Security Centre | Prompt injection is not SQL injection it may be worse | Current online article

Explains why external content processed by generative AI can create security risks.

[Read source 6](#)

Publication note: this guide is produced by AIGAS. References to professional bodies and publications do not imply their endorsement of AIGAS or this guide.