WHITE PAPER 02 / AI AND EVIDENCE

Using AI in IAM Assessments Without Losing Human Judgement

A working review model for generated questions, participant chat, document analysis, sentiment and assessment findings.

AIdentX Editorial | September 2026

Executive summary

AI can help people prepare questions, understand unfamiliar language and examine large amounts of assessment material. Its usefulness depends on what reviewers can verify. A fluent explanation is not evidence that a control works. This paper follows the assessment lifecycle and sets out practical checks for each use of AI, from question generation to the review of risks and gaps.

What you will take away

- Give AI a specific task and a defined scope.
- Separate participant answers, source evidence and generated interpretation.
- Treat sentiment as a prompt for clarification, never as proof of truth.

Written for assessment leaders, reviewers and participants.

The worked scenarios are fictional. The checklists and examples are practical editorial guidance, not customer results or a certification method.

1. Decide what assistance is actually needed

Consider a fictional organisation, Cedar Vale Services, assessing how temporary supplier access is approved and removed. The campaign owner needs relevant questions. A participant needs help understanding what evidence to provide. The approver needs to resolve a difference between an answer and an uploaded procedure. Each task calls for different assistance and a different review.

Define the expected output before invoking AI. A question draft should be specific enough to answer. A document analysis should distinguish what the document says from what it leaves unanswered. A proposed finding should point back to the condition and evidence that justify it. Without this separation, a helpful drafting tool can become an unexamined decision maker.

NIST AI 600-1 describes risks including confabulation and problematic human reliance on generated output. These are reasons to verify important claims rather than equate confidence of expression with correctness. [NIST AI 600-1: Generative Artificial Intelligence Profile].

Create a short local operating rule: what information may be submitted, who checks the result and where the decision is recorded. Make that rule available to participants as well as administrators. The checks below are recommended working practices, not a claim that every deployment enforces them automatically.

2. Review generated questions before launch

In AssessX, the campaign can use questions from the knowledge base and questions generated with AI for the selected scope. The owner should inspect the proposed set before participants receive it. Look for questions that are repeated, too broad, outside the selected application or based on a process the organisation does not use.

For supplier access, a useful generated question might ask how the authorised end date is recorded. A less useful question might assume that a specific identity platform automatically disables every account. Revise it to ask how the outcome is achieved, then request evidence of that outcome. Keep one clear objective per question.

Review check	What to look for
Scope	The question addresses the selected systems and identity population.
Clarity	A participant can understand the request without guessing the terminology.
Evidence	The requested record or example could demonstrate the answer.
Neutral wording	The question does not suggest that a positive answer is expected.

Review check	What to look for
Coverage	Critical stages are included without repeating the same issue.

Where a generated question is retained, the campaign owner remains responsible for its suitability. Saving a poor question into a reusable library can spread the ambiguity into future campaigns. Review reusable content separately from the immediate need to finish this assessment.

3. Help participants explain what actually happens

Participant chat is most useful when it clarifies a question or identifies the kind of evidence being requested. For example: “What does an access expiry control mean in a process that uses tickets?” A useful explanation would describe recording an end date, carrying out removal and checking the result, while allowing for different implementations.

The participant should answer from the organisation’s actual practice. Do not copy a model answer because it sounds complete. If the team records an end date but has no consistent removal check, say so. That limitation is valuable information for an assessment intended to identify improvements.

> **A better participant answer**
>
> “The sponsor includes an end date in the request. The service desk removes access manually. We can provide removal records for two cases, but have not located the third.” This fictional answer separates the process, the available evidence and the remaining uncertainty.

If chat advice conflicts with local policy or system behaviour, ask the campaign owner or reviewer to resolve it. The answer should preserve the participant’s meaning, including uncertainty and exceptions. Good assistance makes the answer clearer; it should not make the control appear stronger than the evidence allows.

4. Keep document analysis attached to its source

Uploading a procedure can help a reviewer understand the intended process. It cannot, by itself, demonstrate that every account followed the procedure. Check the document version, approval status, application coverage and effective date. An analysis of an outdated procedure may be accurate about the file and still irrelevant to the current assessment.

Before upload, use the organisation’s approved data-handling rules. Remove material that the task does not need, such as unrelated personal information or secrets. Check the deployment’s configured processing arrangements rather than assuming that all documents remain in a particular location. For images or scanned pages, verify that important tables and annotations were actually read.

After analysis, compare the important statements with the original document. If a summary says expiry is automatic, find the passage that supports that claim. If there is no supporting passage, treat the statement as unverified and correct the review record. Record contradictions between the file, the answer and operational evidence as follow-up questions.

- Confirm the source, date and scope of the uploaded material.
- Check the passages behind conclusions that affect the assessment.
- Separate a documented requirement from proof of execution.
- Leave missing information visible instead of filling it with an assumption.

5. Use sentiment carefully during approval

AssessX can analyse sentiment in participant answers and expose that information to approvers. The label is an interpretation of language. It is not a lie detector, a measure of staff competence or independent proof that a control succeeds or fails. Short answers, second-language writing and cautious professional wording can all affect how text is interpreted.

An uncertain answer may identify a useful place for follow-up. Ask which part of the process the participant cannot confirm and which record would resolve it. A positive answer still requires evidence. The approver's decision should rest on the answer, supporting material and review rationale, with sentiment used as context.

Automatically generated risks and gaps need the same discipline. Check their source answers, affected scope, wording and potential duplication. A missing removal record may justify a gap in evidence or process assurance. A claim that a supplier definitely accessed sensitive information would need additional evidence. Keep the finding within what the assessment supports.

6. Review the result before it drives action

Use a small set of deliberately difficult examples when evaluating the workflow: a confident answer with weak evidence, a hesitant answer with strong evidence, an outdated policy and contradictory records. Check whether the resulting analysis preserves those distinctions. Test again when the configured model or assessment templates change.

Output	Human decision
Question draft	Is this relevant and answerable?
Chat explanation	Does it help the participant describe the real process?
Document analysis	Are the important statements supported by the source?

Output	Human decision
Sentiment label	Is clarification needed, without judging the person?
Generated finding	Does the evidence justify the condition and consequence?

Measure the effort needed to reach an acceptable result, including review and correction. Counting generated questions or summaries alone rewards volume. A more useful local measure is how often an output can be accepted after review, how often it needs substantive correction and whether unresolved uncertainties are carried into the final decision. These measures help improve the process without claiming that automation replaces accountability.

References and further reading

Primary guidance used for the points cited in this paper. Examples, templates and recommended working practices are AIdentX editorial guidance. Sources reviewed September 2026.

NIST AI 600-1: Generative Artificial Intelligence Profile

https://doi.org/10.6028/NIST.AI.600-1

Put the guidance into practice

Read related articles and watch AssessX product demonstrations at www.aidentx.com/assessx/resources.