

SecurITPosture

Customer guide

Working with the AI

What the AI drafts, the final review, the sign off, and what confidence means.

Version 0.1.0

Generated 2026-10-16

Draft: steps are specified but not final.

Vesta Cyber Solutions. Made from the same source as the guide on the web, so the two always match.

Contents

What the AI drafts	3
Drafts flow, people decide at the end	3
The final review	3
Accepting a part or a whole family	4
The sign off is the one hard stop	4
The disclosure page	4
What confidence means	5

Draft, pending the appliance build. The model that runs inside the appliance, the working view that shows drafts inside the documents, and the wording of the disclosure page are not final. The list of narrative sections the AI drafts is awaiting review.

SecurITPosture includes a small open source AI model that runs inside the appliance. Nothing it reads or writes leaves the virtual machine, and it is never trained on your data. The AI is an assistant. People decide what is final.

The AI is on by default in the appliance. A system administrator can switch it off in Settings, AI assistance.

What the AI drafts

The AI drafts content for the provider's documents. It never writes anything the assessor records or signs: not the SAP, the SAR, the RET, the SRTM, evidence review verdicts, test results, assessor summaries, or any attestation.

- Evidence matching. Each time a customer user or compliance analyst uploads an evidence file, or a new version of one, the AI matches that file against the controls it is likely about. It checks only the new file, in the background, and the upload does not wait for it. Earlier versions are kept.
- Control implementation statements. For controls with matched evidence, the AI drafts the statement and proposes an implementation status. The status is a proposal. The live status changes only when a person accepts the draft. When an integration is connected, the draft includes the resources matched to the control.
- Gap flags. The AI names evidence that is missing or weak for a control, and the evidence categories that would close the gap. It does not write the missing document. A flag creates a task for the compliance analyst, which clears when the evidence is sufficient or a person rejects the flag.
- SSP narratives. The AI drafts narrative sections of the SSP from the questionnaire and system records. When an answer it read changes, it redrafts only the sections that read that answer.
- Score explanation. On the posture report, a customer user or compliance analyst can ask for a plain language explanation of the vulnerability posture score. The product shows it only if every number in it appears on the page; otherwise it shows the derivation alone. The explanation never changes the score.

Where a person has written or accepted a statement, the product keeps that text and shows the newest AI draft beside it as a suggestion. It never overwrites what a person wrote.

Drafts flow, people decide at the end

There are no approval stops while work is under way. Drafts fill the working documents as evidence arrives, marked as AI drafted. Gates and the posture score count only what a person has accepted.

The final review

Open the AI final review, also listed in Settings. Customer users, compliance analysts, and approving officials can open it.

For the selected system, the page lists every AI drafted part, grouped by document and control family. For each part it shows:

- the draft text, marked as AI drafted;
- why the AI wrote it, as plain text;
- the evidence it cites;
- its confidence, with each check explained;
- the gaps the AI could not close.

What each role does:

- A compliance analyst edits a draft or returns it for revision with a note.
- An approving official accepts or rejects drafts.
- While the setting "requester and approver must differ" is on, a person cannot accept a draft they requested or edited.

Accepting a part or a whole family

1. Accept one part at a time, or accept a whole control family or a whole document in one action.
2. Before a bulk acceptance, the product shows how many drafts it covers and how many are low confidence.
3. Each draft is then accepted on its own, with its own audit record. Any draft you may not decide is skipped and listed.

Drafts are never accepted by confidence. A model's confidence can be steered by the text it reads, so bulk acceptance works only by family or document.

The sign off is the one hard stop

The package sign off by the system's accountable official is the one step that blocks. If any AI drafted part included in the package is not accepted, the sign off is refused and the product lists the parts still open. Unaccepted AI text never reaches a signed or submitted package. A draft export from the working view marks every unaccepted AI part.

The disclosure page

When the setting "AI disclosure page" is on, which is the default, each generated provider document includes one page listing every AI assisted section and who accepted it. The package metadata records AI assistance whatever the setting says.

What confidence means

Confidence is not the model's opinion of itself. The product computes it from checks it can verify, and shows each one in the final review:

- Evidence support: whether the evidence the draft cites exists, and whether its statements are backed by text in that evidence. Sentences without support are marked.
- Requirement coverage: how much of the control's requirement, including its lettered parts and parameters, the statement addresses.
- Form: whether the draft is well formed, the right length, and free of stray labels or instructions.
- The model's own rating, which counts the least.

The product then calibrates the result on a set of examples with known answers, so that 0.80 means about 80 percent of such drafts are correct.

- "Uncalibrated" means no calibration exists yet for the current method. The number is then a ranking, not a chance of being correct. Drafts below the threshold are still marked for review.
- "Low certainty" means too few examples, or too much error, stand behind that score. The page shows how many examples there were.

Drafts below the review threshold (80 percent by default) are marked low confidence. They still appear in the working documents and wait for the final review like any other draft.