

**Violet Attention Audit™**

Constructed demonstration - not client work or an observed result

WHAT VIOLET EXAMINED

AI-assisted application review announcement

Submitted message

"We are introducing AI to review applications more quickly and consistently. This will improve processing times and help us deliver better service."

RECEIVER

Applicants and prospective applicants

RECEIVER'S JOB

Understand what AI does, who decides, and what happens after a decline

INTENDED TRANSITION

Exposure to trusted, decision-ready understanding

REPORT USER'S DECISION

Decide whether the announcement is ready to publish

AUDIT VERDICT**NOT YET SENDABLE**

Revise before publication. The message promotes efficiency before explaining decision authority, accountability, and recourse.

Publication condition: Confirm that a trained reviewer makes the final eligibility decision and every declined applicant can request human reconsideration. If either condition is false, the message and process must change together.

RECOMMENDED REDESIGN**CUSTOMER ANNOUNCEMENT****How AI will support application review**

We are adding AI to help our review team organize application information and identify details that may need closer review.

AI will support the review; it will not make the final eligibility decision.

1

What remains accountable to a person

A trained reviewer will assess the application, make the final decision, and remain responsible for checking the information used.

2

If your application is declined

Your decision notice will explain the reasons and include a link to request human reconsideration. A person who was not responsible for the original decision will review that request.

3

You do not need to take action now. We will explain any application-process changes before they take effect.

Next action

Confirm final-decision authority and the reconsideration process with the accountable product, operations, and legal owners; then replace the submitted announcement with the verified version above.



VIOLET 5S™ DIAGNOSIS

Primary break: Substance + Stake

What already works: The original announcement states that a change is coming and names the intended operational benefit: faster, more consistent review. Those useful facts should remain after the authority and recourse questions are answered.

Spotlite

Efficiency receives attention before the applicant's decision risk.

Substance

The message does not define what AI does or does not decide.

Source

The accountable human decision-maker is unnamed.

Strain

Applicants must infer authority and search elsewhere for recourse.

Stake

A decline makes authority, accountability, and challenge rights salient.

Priority: Substance + Stake are the primary break; Spotlite supports it; Source is an operational dependency. The likely break is exposure to trusted, decision-ready understanding - not exposure itself.

1

Define the AI boundary

Original problem: "AI will review" blurs assistance and final authority.

Correction: Name the support task and the final-decision limit.

Check: Can applicants explain AI's role without seeing the wording again?

2

Name the accountable person

Original problem: No person is named as responsible for the decision.

Correction: Put human responsibility immediately after the AI boundary.

Condition: Do not use this claim unless the operating process supports it.

3

Place recourse beside the risk

Original problem: The submission gives no path after an adverse decision.

Correction: Put human reconsideration beside the decline risk.

Check: Verify eligibility, independence, timing, and the destination link.

Confirm before implementation

- Does a trained reviewer always make the final eligibility decision?
- Can every declined applicant request human reconsideration, and is the second reviewer independent of the original decision?
- Which privacy, fairness, retention, or regulatory facts must also appear in this announcement or the connected process?

Verify that the change helped**Primary**

Customers correctly explain the AI's role, final authority, and reconsideration route.

Compare

Test the current and proposed announcements under the same task and context.

Guardrails

Track unsupported reassurance, new confusion, complaints, and whether the process actually honours the message.

Weaken

Revise the diagnosis if customers already understand these points or if distrust remains because the underlying policy is unacceptable.

Evidence boundary: This is a constructed public demonstration. Violet did not examine a client's complete policy, interview customers, observe the journey, or measure an outcome. The receiver questions and predicted improvements are hypotheses to test.

Guardrail: Clearer communication cannot repair an unfair, inaccessible, or unaccountable review process. Preserve informed choice and the ability to challenge a decision.

Method note: No visual-priority model or eye-tracking evidence was used. The layout hierarchy is a design proposal, not a measurement of attention.