PRACTICE 1 OF 4 · HUMAN-OWNED

# Name the release owner before the first prompt

*Who has responsibility and authority to decide this is ready to use?*

Fill this in before anyone opens the AI tool, and put it where the team will see it: the project plan, the top of the project folder, or the intake form. Ask "who decides?" while it is cheap to answer, not at 4:30 on the Friday before go-live.

## The statement

| | |
|---|---|
| **Project** | |
| **Release owner (a person, not a role or a team)** | |
| **Evidence they will decide on** | |
| **Decision due by** | |

**[Name] decides this is ready to use, based on [evidence], by [date]. If [name] can't, it doesn't ship.**

## The decisions behind it

"Is the content accurate?" and "Is it good training?" are different decisions, often owned by different people. Someone then has to weigh both and say "use it."

| Decision | Named person | Evidence they need | By when |
|---|---|---|---|
| **Content accuracy** | | | |
| **Instructional quality** | | | |
| **Release (weighs both)** | | | |

## Check before you start

- It names a person. "The SME" or "L&D" is a role, and a role can't decide anything.
- It names the evidence. That sets up the source register (practice 2) and the claim review (practice 3).
- It has a default: if the owner can't decide, it doesn't ship.
- It was written down before the first prompt.


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

**WORKED EXAMPLE** · The fictional lost-laptop project from the session. Names, extensions and dates are invented.

| | |
|---|---|
| **Project** | Lost or stolen company laptop: 10-minute module and one-page job aid |
| **Release owner** | Dana Ruiz, L&D lead |
| **Evidence they will decide on** | Security Operations' claim-by-claim sign-off on v0.4, plus a two-person user walkthrough of the job aid |
| **Decision due by** | Friday, 12 March |

| Decision | Named person | Evidence they need | By when |
|---|---|---|---|
| **Content accuracy** | Sam Okafor, Security Operations manager (owns the procedure) | Every claim in v0.4 checked against Rev. C | Wed 10 Mar |
| **Instructional quality** | Dana Ruiz, L&D lead | Walkthrough notes; quiz items aligned to objectives | Thu 11 Mar |
| **Release (weighs both)** | Dana Ruiz, with Sam's sign-off in hand | Both of the above, filed in the release record | Fri 12 Mar |

*Dana Ruiz decides this is ready to use, based on Security Operations' claim-by-claim sign-off and a user walkthrough, by Friday, 12 March. If Dana can't, it doesn't ship.*


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

PRACTICE 2 OF 4 · SOURCE-GROUNDED

# Keep a source register, and make conflicts visible

*What authoritative information should this output reflect?*

Before the first prompt, list everything that might go into the AI tool and what each source is for. Three statuses do most of the work:

- **Authoritative.** The output must match it. Every claim is checked against it.
- **Reference.** It can shape structure or tone, but it doesn't get the final word on facts.
- **Excluded.** It doesn't go in at all. Say why, so nobody uploads it later "for context."

## The register

| Source | Status | Version · owner | How it's used |
|---|---|---|---|
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

## The prompt line

Add this to every prompt that uses more than one source:

**"If the sources disagree, list the disagreement. Don't resolve it."**

The AI is good at noticing that two documents say different things. It has no authority to decide which one is right. That decision belongs to the owner of the authoritative source.

## Conflict log

| What disagrees | Sources involved | Sent to | Resolution and date |
|---|---|---|---|
| | | | |
| | | | |
| | | | |


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

**WORKED EXAMPLE** · The fictional lost-laptop project from the session. Names, extensions and dates are invented.

| Source | Status | Version · owner | How it's used |
|---|---|---|---|
| **Lost or stolen device procedure** | Authoritative | Rev. C, effective 1 March · Security Operations | Every claim must match it |
| **Last year's module** | Reference | v2.1 · L&D | Structure and tone only |
| **2023 lost-device FAQ** | Excluded | Superseded | Not uploaded: it still says 24 hours to the Service Desk |

| What disagrees | Sources involved | Sent to | Resolution and date |
|---|---|---|---|
| Reporting window and route: 1 hour to Security Operations vs. 24 hours to the Service Desk | Rev. C vs. last year's module | Sam Okafor, Security Operations | Rev. C governs. Module text and quiz item 4 updated. 3 Mar |

Why this matters: in the session's example, the old 24-hour rule survived in quiz item 4, and the feedback text credited it to Rev. C. The citation pointed to the right document; the claim came from the wrong one.


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

PRACTICE 3 OF 4 · VERIFIABLE

# Check claims, not screenshots

*What checks would establish that it is fit for its intended use?*

Pull every actionable claim out of the draft: steps, deadlines, phone numbers, who to contact, anything someone would act on. Put each one next to the line in the authoritative source that supports it. If you can't find a source line, that's a finding.

| | |
|---|---|
| **Product and version reviewed** | |
| **Authoritative source and version** | |
| **Reviewer** | |
| **Sent on / due by** | |

## The review request

**"Please review [version], items [1–n], against [source]. For each one: approve, change, or can't tell."**

"Can't tell" matters. It lets the reviewer say the source itself is unclear, which is far more useful than a polite "looks good."

| # | Claim in the draft | Source location | Response | Notes |
|---|---|---|---|---|
| | | | | |
| | | | | |
| | | | | |
| | | | | |

## User walkthrough (ten minutes, if you can)

Have two people from the target audience use the product on a realistic scenario. You'll learn things no review will tell you.

| Who | Scenario | Where they hesitated or went wrong |
|---|---|---|
| | | |
| | | |


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

**WORKED EXAMPLE** · The fictional lost-laptop project from the session. Names, extensions and dates are invented.

| | |
|---|---|
| **Product and version reviewed** | Lost-laptop job aid, v0.4 |
| **Authoritative source and version** | Lost or stolen device procedure, Rev. C |
| **Reviewer** | Sam Okafor, Security Operations |
| **Sent on / due by** | Sent Mon 8 Mar / due Wed 10 Mar |

*"Please review v0.4, items 1–9, against Rev. C. For each one: approve, change, or can't tell."*

| # | Claim in the job aid (v0.4) | Source location | Response | Notes |
|---|---|---|---|---|
| 1 | Report within 1 hour | Rev. C, §3.1 | Approve | |
| 2 | Call Security Operations at [ext.] | Rev. C, §3.2 | Change | Use the 24/7 line, not the office extension |
| 3 | Then file a Service Desk ticket | Rev. C, §3.4 | Approve | |
| 4 | Report even if you think you just misplaced it | Rev. C, §3.1 | Approve | Missing from v0.4; restored in v0.5 |
| 5 | Change your passwords from another device | None found | Can't tell | Not in Rev. C. Sent to Security to decide |

| Who | Scenario | Where they hesitated or went wrong |
|---|---|---|
| Two field staff | "Your laptop isn't in your bag at the airport gate." | Both looked for the phone number first; it was at the bottom. Moved to the top in v0.5. |


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

PRACTICE 4 OF 4 · TRACEABLE

# Release record

*What record would let someone understand how this version was produced and approved?*

One page, filed next to the product, written by a person at release. Most of it comes from the other three practices. If it takes more than ten minutes to fill out, it's too long.

| | |
|---|---|
| **Product** | |
| **Version and date released** | |
| **Release owner** | |
| **Sources used, with versions** | |
| **AI tool and what it produced** | |
| **Human changes made after the draft** | |
| **Checks: who, what, result** | |
| **Approval: who, which version, when** | |
| **Review trigger** | |
| **Known limits or open questions** | |

The review trigger tells future you when this product stops being trustworthy, for example "any revision to the procedure." A chat transcript is the AI's account of the work. This page is the record of the decisions.


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.

**WORKED EXAMPLE** · The fictional lost-laptop project from the session. Names, extensions and dates are invented.

| | |
|---|---|
| **Product** | Lost or stolen company laptop: 10-minute module and one-page job aid |
| **Version and date released** | v1.0, released 12 March |
| **Release owner** | Dana Ruiz, L&D lead |
| **Sources used, with versions** | Procedure Rev. C (authoritative); last year's module v2.1 (reference, structure only). 2023 FAQ excluded. See the source register. |
| **AI tool and what it produced** | Organization's approved generative AI tool: first draft of module script, job aid and 8 quiz items (v0.1) |
| **Human changes made after the draft** | Restored the "report even if misplaced" step the AI dropped; corrected quiz item 4 to 1 hour; moved the 24/7 number to the top of the job aid |
| **Checks: who, what, result** | Sam Okafor: claim review of v0.4, 9 items against Rev. C, 1 change, 1 can't tell (resolved). Two-person walkthrough: 1 layout change. |
| **Approval: who, which version, when** | Sam Okafor (accuracy) and Dana Ruiz (release), v1.0, 12 March |
| **Review trigger** | Any revision to the lost or stolen device procedure |
| **Known limits or open questions** | Password-change advice left out pending a decision from Security |


Free to copy and adapt for your own work, with credit. Licensed CC BY 4.0.