


FIELD BRIEFING

A better brief. A better answer.

Context is part of the work.



What should I give an AI before I ask it to help?

The instruction “make this better” leaves several decisions unresolved. Better for whom? For which purpose? Using which facts? A useful brief gives the model enough context to work within the task and gives the reviewer a way to judge the result.

Explanatory schematics and worked examples are illustrative. No live model, observed client results or measured market performance are presented.

Read the mechanism.

Five parts make a practical starting point: the goal, the audience, the source material, the constraints and the expected output. This is a KAF writing framework, not a special syntax that every model requires. Use only the parts the task needs.

For an invented supplier update, the goal is to explain a delivery delay; the audience is a customer waiting for an order; the source says the new delivery window is Tuesday to Thursday, with no exact day confirmed. The constraint is not to invent compensation or a promise. The output is a short email with a proposed subject line and a separate list of facts requiring confirmation.

A vague request might produce a confident promise about Tuesday morning. A grounded brief makes that promise visibly outside the evidence. It does not guarantee the model will obey. The person checking the email must still compare dates, obligations and tone with the source.

Examples help when form matters. Show one approved response with the correct level of detail, then say what it illustrates. Avoid including personal information just to demonstrate style. Clearly separate instructions from documents: a quoted customer email is evidence to consider, not authority to change the task.

Iteration should have a reason. If the output is too long, specify a useful limit. If it makes unsupported claims, state which source material is authoritative and how to show gaps. If it misunderstands an audience, explain what that audience already knows. Adding more adjectives is usually less helpful than naming the observed failure.

Save the brief and model configuration alongside sample results when the workflow matters. A successful one-off answer is a useful example, but repeated performance on representative inputs is stronger evidence.

Technical references: Microsoft · Prompt Engineering; Anthropic · Building Evals; OWASP · Prompt Injection (2026). Links and source scope appear in the final section.

Try the idea.

INVENTED TEACHING EXAMPLE

A supplier update grounded in known facts

Goal: explain a delay. Source: delivery expected Tuesday–Thursday; no exact day confirmed. Constraints: no promised compensation or exact date. Output: subject line, short email and confirmation checklist. Illustrative response: “Your delivery is now expected between Tuesday and Thursday. We do not yet have a confirmed day.” This is prepared text, not a live model result.

Go deeper.

Separate evidence and instructions

Delimit supplied documents and state what role they play. This helps readers and models understand the task, but text delimiters are not a security boundary. Tool permissions and validation still apply.

Version a useful brief

Record the prompt text, model/version and relevant settings when the task is repeatable. Keep development examples separate from evaluation examples. A saved success is an example, not a reliability estimate.

Test unknowns

Include inputs where dates, amounts or ownership are absent. Require an explicit unknown field or confirmation list, then check whether the model actually avoids inventing the missing value.

Keep the evidence close.

Constraint

A limit or requirement the output should respect.

Source material

Information provided as evidence for the task, not instructions to change it.

Example

An illustration of the required form or behaviour, not proof of general performance.

Acceptance rule

A specific condition used to judge the result.

Sources · accessed 4 October 2026

[Microsoft · Prompt Engineering](#)

[Anthropic · Building Evals](#)

[OWASP · Prompt Injection \(2026\)](#)

Technical and security references support the mechanisms, not our fictional results. KAF's worksheets and examples are original editorial recommendations. No current provider pricing, comparative ranking or legal compliance is asserted. Verify the exact release, configuration and policy for your use.

Practical AI Insights · First editorial edition · 4 October 2026

A better brief. A better answer. · [Web edition](#)