The ENABLE Prompt Framework (v6)

November 6, 2026 · Richard Turner · aligned to TURNS v13.2 and book v13

Making AI human means people stay in charge of purpose, evidence, and judgment. ENABLE is Thinking Enabled's six-step framework for writing clear AI requests and checking what comes back: Establish the aim, Name the perspective, Anchor in context, Build the brief, Limit the scope, then Evaluate and take another turn.

Quick reference card

	Step	What you decide	Ask yourself
E	Establish the aim	The purpose, the people served, and what success looks like	What will this help someone do?
N	Name the perspective	The expertise or viewpoint that sharpens the answer	Whose judgment should shape this?
A	Anchor in context	Background, sources, examples, and which facts are confirmed or assumed	What must the AI know?
B	Build the brief	The deliverable, format, method, and checkpoints	What should I receive, and where should it pause?
L	Limit the scope	Audience, tone, length, exclusions, privacy, and human approval	What must it respect, avoid, or leave to me?
E	Evaluate	The checks, feedback, and revisions	What will I verify, and what should change?

After Evaluate, take another turn: revise the prompt or the work, then run it again.

How much ENABLE does this task need?

Match the effort to the cost of being wrong.

Tier	Use it when	Steps	Example
Quick	The stakes are low and the result is easy to check	Establish, Build, Limit	Tightening a short email
Standard	Most everyday work	All six	A client update or a one-page brief
High-stakes	Others will act on the result, or it touches money, people, or public claims	All six, plus a named approver, checkpoints, and	A pilot plan, a policy summary, an analysis for

Tier	Use it when	Steps	Example
		independent verification	a decision

The six steps in practice

Establish the aim. Start with the reason, not the request. "Write a client update" is a task. "Write a client update that explains a missed deadline and requests approval for a new date" is an aim. State success in terms you can observe.

Name the perspective. Choose a lens that changes what gets noticed: a skeptical CFO, a patient workshop facilitator, a first-time customer. Say whether you want drafting, analysis, critique, or recommendations. A role guides emphasis; it does not give the AI credentials.

Anchor in context. Supply the notes, data, policies, and examples the answer should rest on. Separate confirmed facts from assumptions, and ask the AI to flag missing or conflicting evidence. A rule against invention is no substitute for good source material.

Build the brief. Name what you want in your hands: an email, a two-page brief, a comparison table, working code. Add the required elements and the number of versions. Ask for clarifying questions, or an outline and approval, when a wrong guess would be costly.

Limit the scope. Set audience, tone, length, and exclusions. State privacy rules and where human approval is required. For factual work, require supported claims and visible assumptions. For fiction, keep invented content apart from claims about real people.

Evaluate. Check the output against your aim and acceptance criteria. Verify consequential facts at the source and recompute calculations. Give specific feedback, such as "Lead with the revised deadline and shorten paragraph two," then take another turn. Self-review helps a draft, but it is not independent verification.

Fill-in templates

Replace each [PLACEHOLDER] and delete any line you do not need.

Compact (Quick and Standard tasks)

```
Help me [TASK] so that [AIM], using [CONTEXT AND SOURCES]. Take the view of
[PERSPECTIVE] if it helps. Return [DELIVERABLE] for [AUDIENCE], within
[LIMITS]. Flag evidence gaps and check the result against [SUCCESS CRITERIA].
```

Full (Standard and High-stakes tasks)

```
ESTABLISH: Help me [TASK] in [SITUATION] so that [PERSON OR TEAM] can
[DECISION OR ACTION]. Success means [OBSERVABLE CRITERIA].

NAME: Take the perspective of [EXPERTISE OR STAKEHOLDER]. Your job is to
[DRAFT, ANALYZE, CRITIQUE, OR RECOMMEND], considering [PEOPLE AFFECTED].

ANCHOR: Use [BACKGROUND] and [SOURCE MATERIAL]. Treat [CONFIRMED FACTS] as
```

```
established and [ASSUMPTIONS] as unverified. Flag missing or conflicting
evidence. For facts that change, use [SOURCE AND DATE REQUIREMENTS] or say
what still needs checking.

BUILD: Produce [DELIVERABLE AND FORMAT] containing [REQUIRED ELEMENTS]. Use
[METHOD OR SEQUENCE] if it applies. Ask only questions that could materially
change the result; otherwise state your assumptions. Pause for my approval
before [CHECKPOINT].

LIMIT: Write for [AUDIENCE] in [TONE AND STYLE], within [LENGTH AND SCOPE].
Respect [PRIVACY AND ACCESSIBILITY NEEDS]. Exclude [EXCLUSIONS]. Do not
present invented facts, figures, quotes, or sources as real. Label
assumptions. Leave [RESERVED DECISIONS] to [HUMAN APPROVER].

EVALUATE: Check your response against [ACCEPTANCE CRITERIA] and the supplied
evidence. List unresolved issues and any claims I should verify
independently. After my review, revise according to [SPECIFIC FEEDBACK].
```

Before and after

A clear request removes avoidable guesswork and makes the response easier to judge. All details below are fictional and for teaching.

Example 1: a low-stakes email, three levels

Before

> *Write an email about our workshop.*

Good enough (Quick)

> *Help me write a short email inviting a prospective COO to a 20-minute conversation about a pilot AI literacy workshop. Keep the body under 150 words, warm and plain, with no invented client results. I will review it before sending.*

Full (Standard)

> *ESTABLISH: I run AI literacy workshops through Thinking Enabled. Help me invite a prospective COO to discuss whether a pilot workshop would suit the team. Success is a clear, relevant invitation with one manageable next step.*
>
> *NAME: Take the perspective of a B2B consultant who writes plain, warm outreach and respects the reader's time.*
>
> *ANCHOR: Use this approved description: "A 3.5-hour facilitated workshop in which leaders choose one AI pilot candidate worth investigating and plan how to check it with the people who do the work." The COO's priorities are unconfirmed, and we have no approved client results to cite.*
>
> *BUILD: Write one email with a subject line and two alternates. Introduce the workshop, explain its relevance without claiming to know the COO's circumstances, and invite a 20-*

> *minute conversation. Draft directly unless missing information would change the message.*
>
> *LIMIT: Body under 150 words, confident and approachable. Do not invent results, statistics, personalization, availability, or guarantees. Use a name placeholder. I will review and send it myself.*
>
> *EVALUATE: Check that the email matches the description, has one next step, and makes no unsupported claims. Report the body word count separately.*

What to check. The workshop description is accurate, the email does not invent the COO's priorities, the body stays under 150 words, and the subject lines promise nothing the workshop cannot deliver. A useful follow-up: "The opening is too general. Lead with the chance to choose one pilot candidate, keep the 20-minute invitation, and cut the body to 120 words. Do not add benefit claims."

Example 2: a high-stakes analysis

Here Anchor and Evaluate carry most of the weight.

> *ESTABLISH: A 40-person team piloted an AI drafting tool for six weeks. Help me summarize the pilot for a director deciding whether to expand it. Success means the director can see what was measured, what was estimated, and what is unknown.*
>
> *NAME: Take the perspective of a skeptical operations analyst who separates evidence from opinion.*
>
> *ANCHOR: Use the attached time-tracking export and survey summary. Treat the time-tracking figures as measured and the managers' quality ratings as opinion. Flag any conflict between sources.*
>
> *BUILD: Produce a one-page summary with findings, strength of evidence, risks, and the four Renewal calls: continue, adjust, expand, or stop. If the pre-pilot baseline is unclear, ask me before drafting.*
>
> *LIMIT: Plain language for a non-technical director. Do not recommend expansion. Label any threshold as provisional until measured and approved. The director is the Decider and makes the call.*
>
> *EVALUATE: Recompute every figure from the export. List the claims I must verify independently, and mark anything unmeasured as "not yet measured."*

Tune a library prompt, and save what works

Prompts from a prompt library usually arrive with the aim, the brief, and the limits already written. They rarely carry your real context or your own acceptance check. Three moves close the gap:

1. **Anchor.** Replace each placeholder with real material, and mark what is confirmed and what is assumed.
2. **Name.** Add a perspective if the prompt has none and the task would benefit from one.

3. **Evaluate.** Add the success criteria and the claims you will verify yourself.

For example, a library-style prompt such as "Draft a status update for [AUDIENCE] using [PROGRESS NOTES]" becomes stronger when you add: "Treat the attached notes as the only source. Flag anything that looks out of date. List any claim I should confirm with the project owner before I send it."

Save what works. When a prompt earns its keep, keep it as a reusable prompt:

- Record what you changed from the original and why.
- Run it on a fresh case before you trust it again, and note any failures.
- Keep practice examples separate from the cases you use to judge it.

This is the same habit as the last TURNS stage: write down what you learned so the next turn starts further ahead.

How ENABLE fits inside TURNS

TURNS guides the broader work of adopting and using AI. ENABLE frames and improves each request inside that work. Each TURNS stage borrows a habit from Richard's four worlds: UX (start with the humans), AI (prototype fast, on real data), TV (rehearse before you go live), and DOD (clear intent, rules of engagement, after-action review).

TURNS stage	Habit it borrows	ENABLE steps it feeds	What it adds
T: Target the Mission	DOD commander's intent	Establish, Build	A Target Statement (purpose, end state, evidence, limits) with an owner and a Decider
U: Understand the Humans	UX research	Name, Anchor, Limit	Real workflows, pain points, and a trust baseline; leader hypotheses marked unverified
R: Rehearse with AI	TV rehearsal, AI prototyping	Anchor, Build, Evaluate	Approved real examples, compared with a baseline; practice cases kept apart from final tests
N: Navigate the Rules	DOD rules of engagement	Build, Limit	Human review points, approval authority, and off-limits data and actions
S: Sustain and Debrief	DOD after-action review, TV post-production	Evaluate, Anchor	A recorded gate call, updated evidence, a review rhythm and, after a go, the 20-weekday plan (Daily Milestone)

Use the same judgment at both levels. TURNS sets the context and the accountability. ENABLE carries that context into the request. A well-written prompt does not authorize an action, establish a baseline, validate an assumption, or approve expansion. A pilot-planning prompt should name the human approver, separate measured results from estimates, and label proposed thresholds as provisional until measured and approved.

In The Turning Lab and the courses. Ask participants to pick a real task, name its TURNS stage, and draft the request with ENABLE. Have them check one response against the original evidence and criteria, then revise the prompt or the work. Groups without an approved AI account can check a supplied output instead.

Use ENABLE with book v13

TURNS governs the broader work. ENABLE carries the Target Statement, approved source material, human decision rights and test criteria into each request. Keep the six ENABLE steps unchanged.

Label every claim. Kind of claim: Evidence / Interpretation / Validated. Status: Measured / Passed / Failed / Not yet measured / Not applicable with a reason. Record source, date and sample size. A failure is a result, not a gap. Missing evidence is never a pass or a zero.

Write the rehearsal script. In R, the six-part script has Role, Task, Materials, Format, Rules and Stop rule. Name the perspective shapes the Role. Establish the aim and Anchor in context shape the Task and Materials. Build the brief shapes the Format. Limit the scope shapes the Rules and the Stop rule. Evaluate becomes the scoring plan, and the people who know the work still do the scoring. ENABLE helps write and improve the script; it does not replace independent scoring, sealed final cases, a task check or a human gate decision.

Use the gate calls. Greenlight: proceed, revise or park. On Air: go, adjust or no-go. Renewal: continue, adjust, expand or stop. Evidence opens a gate; the calendar does not.

Bayside stays at Draft. The AI prepares; a dispatcher reviews and sends. A prompt never raises the autonomy rung.

Check before you rely

Before you send the prompt

- ☐ **Aim:** I can say who the work serves and how I will use the response.
- ☐ **Perspective:** I chose a useful lens, or left it out because the task is clear.
- ☐ **Context:** I supplied the evidence and separated facts from assumptions.
- ☐ **Brief:** I named the deliverable and any method or checkpoint that matters.
- ☐ **Limits:** I set audience, scope, boundaries, and human approval rules.
- ☐ **Privacy:** I replaced every placeholder and removed information I am not authorized to share.
- ☐ **Evaluation:** I know what success looks like and what I will verify.

Common mistakes and quick fixes

Mistake	What happens	Quick fix
A task with no intended use	Generic, safe output	State the aim and one observable success criterion
A role standing in for	Confident answers with	Provide sources, and say how to handle gaps

Mistake	What happens	Quick fix
evidence	no support	
An unclear deliverable	The wrong format or length	Name the format, required elements, and reader
Questions or approvals at every step	A slow, tiring exchange	Add a checkpoint only where the stakes justify it
Trusting a rule against invention	Errors still slip through	Check consequential claims against the original evidence
Vague feedback such as "make it better"	Random changes	Name what missed the aim and request one specific change
Treating fluent output as verified	Polished mistakes reach other people	Check evidence, calculations, and suitability before use

Making AI Human

Three habits carry the framework:

- **Lead with purpose.** Say what the work should help someone accomplish.
- **Ground the request.** Supply context and evidence, and make uncertainty visible.
- **Review before relying.** Verify what matters, then take another turn.