

THE PROMPT DISCIPLINE

Card Pack

"The discipline behind the deploy."

Five printable one-page cards — one per lever — for teams building Copilot Studio agents, Azure OpenAI integrations, and Power Platform automations that have to hold up in production.

Each card carries the memorable rule, the three components, a five-point pass check, and space to write your own answer for the prompt in front of you.

CARD 01 **Intent** — the first sentence that decides what your prompt costs.

CARD 02 **Context** — the subtractive discipline.

CARD 03 **Boundaries** — the shape a downstream system can consume.

CARD 04 **Examples** — the hardest case, not the easiest.

CARD 05 **Verification** — what good looks like, and how it will be checked.

BY CARMEN NEL

Power Platform & AI Solutions Architect

Intent.

The first sentence that decides what your prompt costs.

THE TEST

“Can I explain what this prompt must do to a three-year-old?”

Not the technology — the task. Said once, in plain words a child could act on. If the instruction needs jargon and three clauses before it makes sense, the intent is unclear to you, not to the model.

THE THREE COMPONENTS

Task	One verb. Summarise, classify, draft, extract. Not 'help' or 'handle'.
Output type	The shape of a correct answer. Paragraph, JSON array, single number, list of four.
Audience	Who reads the output and what they already know. A finance director is not a new joiner.

PASS CHECK

- [] 1. The task is named in one verb.
- [] 2. The output type is stated (format, length, structure).
- [] 3. The audience is named.
- [] 4. Read it aloud — a three-year-old could describe what 'done' looks like.
- [] 5. If any of the four fails, the intent is not ready. Do not paper over it with context.

Context.

The subtractive discipline.

THE TEST

“If I removed half of this context, would the answer still be correct?”

Not ‘could a smaller context work in general’ — a yes-or-no on the actual passage in the actual prompt. Delete the second half, send it, check. If the answer still holds, the second half was noise you were paying for.

THE THREE COMPONENTS

Source	One document or index. If more than one, name which wins. Ambiguity in source is ambiguity in answer.
Scope	The smallest excerpt that still answers. Not the largest excerpt that fits. Retrieval is subtractive too.
Fallback	What the model says when the context does not contain the answer. Silence is a decision. Write it explicitly.

PASS CHECK

- [] 1. The source is named. One index, or one document.
- [] 2. The scope is narrowed — hand-picked excerpt or targeted retrieval.
- [] 3. The fallback is written into the prompt.
- [] 4. Run the half test on the actual passage.
- [] 5. Keep only what survives. Delete the rest.

Boundaries.

The shape a downstream system can consume.

THE TEST

“Could a downstream system consume this output without writing a custom parser?”

Boundaries turn a model from a writer into a component. State the schema, the length, the tone, and what must be excluded. If Power Automate has to defensively parse it, the boundaries are missing.

THE THREE COMPONENTS

Format	The exact output shape. JSON schema, table columns, sentence count. Machine-parseable, not prose.
Length	Explicit bounds. 'At most five items', 'exactly three paragraphs'. 'Concise' is not a boundary.
Exclusion	What must not appear. Preamble, apologies, disclaimers, extra fields. State it, do not assume it.

PASS CHECK

- ☐ 1. The output format is explicit and machine-parseable.
- ☐ 2. Length is bounded with a number, not an adjective.
- ☐ 3. What to exclude is stated, not assumed.
- ☐ 4. A downstream flow can consume the output as-is.
- ☐ 5. The AFTER prompt survives being run twice and produces the same shape both times.

→ *Deep-dive article coming soon on carmennel.co.za.*

Examples.

The hardest case, not the easiest.

THE TEST

“Did I include the example that captures the case I keep correcting by hand?”

A short, well-chosen example does more work than three paragraphs of instructions. The discipline is to pick the example that demonstrates the hardest case — not the tidiest one.

THE THREE COMPONENTS

Anchor	One example that fixes the shape and tone the model must copy.
Edge	The example that captures the case you keep correcting. Not the easy one.
Contrast	Where useful, a negative example: what a wrong answer would look like, so the model can avoid it.

PASS CHECK

- [] 1. At least one example is present.
- [] 2. The example demonstrates the hardest case, not the easiest.
- [] 3. Tone is set by example, not by adjective ('professional', 'friendly').
- [] 4. Where the model gets it wrong, a contrast example is added.
- [] 5. The example matches the intent, source, and boundaries of the prompt.

→ *Deep-dive article coming soon on carmennel.co.za.*

Verification.

What good looks like, and how it will be checked.

THE TEST

“If the output is wrong in a subtle way, does the prompt give us a chance to catch it?”

A disciplined prompt closes with a verification clause: the success criteria, and where appropriate an instruction for the model to self-check before returning. The cheapest evaluation layer available.

THE THREE COMPONENTS

Success criteria	What a correct output must contain. Named fields, arithmetic checks, presence of citations.
Self-check	An instruction for the model to verify the criteria before returning. A pre-flight check.
Failure mode	What the model should return when the check fails. 'verification_failed' with a reason, not partial data.

PASS CHECK

- [] 1. Success criteria are explicit and testable.
- [] 2. The model is asked to self-check before returning.
- [] 3. A failure mode is defined for when the check fails.
- [] 4. The failure mode is machine-parseable by the downstream flow.
- [] 5. The check is cheap enough to run on every call.

→ *Deep-dive article coming soon on carmennel.co.za.*