


The Glass of Water

Imagine a **glass of water**.

Every message you write, every answer the AI gives — that's water being poured into the glass.

When the glass is full, **older content gets cut off**. The AI simply can't see it anymore.

This is the **context window** — the AI's short-term memory. It has a fixed size. Everything beyond it is gone.




Short conversation

Plenty of room



Getting longer

Starting to fill up



Full

AI starts forgetting

How AI Reads

AI doesn't read words like we do. It breaks everything into small chunks called **tokens** — roughly 1–3 per word, depending on the language.

Tokens are the **unit of measurement for the water in the glass**. Every message in, every answer out — all measured in tokens.

That's why a very long conversation gets worse over time — the glass fills up and older context gets cut off.

New conversation = empty glass.

That's why starting fresh often works better than a very long chat.

Same sentence, different languages

ENGLISH

How can I help you?

~ 5 tokens

DEUTSCH

Wie kann ich Ihnen

helfen?

~ 7 tokens

German uses **more tokens** than English for the same content — the glass fills up faster.

The glass needs to be filled

Better models, better prompts — both matter less than **what you put into the glass.**

Empty glass → **generic answer.** Without context — who you are, what you do, how you write — every reply stays standard.

Context compounds. A good prompt helps once. Good context makes *every* next session better.


The skill of 2026 isn't prompting.
It's context management.

IT ALREADY LIVES ON YOUR MACHINE

Claude Code & Markdown

Claude Code reads local Markdown files the moment it starts — your rules, your projects, your voice.

Your context is a **file, not a product.** No cloud, no lock-in. Yours to version, yours to share.

 `/your-project/CLAUDE.md`

What changed in the last year?

One fundamental thing: AI went from **answering questions** to **doing work**.

BEFORE

AI answers

You ask, it responds. One question, one answer. Like a very smart search engine.



NOW

AI works

You give it a task, it plans the steps, executes them, checks the result. Like a capable assistant.

This is the shift from **chatbot** to **agent**. And it changes everything about what a single person can accomplish.

But what *is* an agent?

CHATBOT

You ask



AI answers



You ask again



AI answers

You drive. AI is the **navigation**.

AGENT

You set a goal



AUTONOMOUS
LOOP

▶ Plan next step



🔧 Use tools



✓ Check result

repeat ——— ↑

AI becomes the **driver**.

What made this real

The idea of AI agents isn't new. What's new: **they're now connected to your computer.**



The Command Line

AI no longer lives in a chat window. It runs **directly on your machine.**

Reads your files. Creates new ones. Runs commands. Tests its own work. Like a developer sitting at your keyboard.



A Model That Can Code

The AI behind it is **exceptionally good at programming** — better than most humans.

Claude Opus 4 by Anthropic. Understands complex systems, writes production-ready code, and reasons through problems step by step.



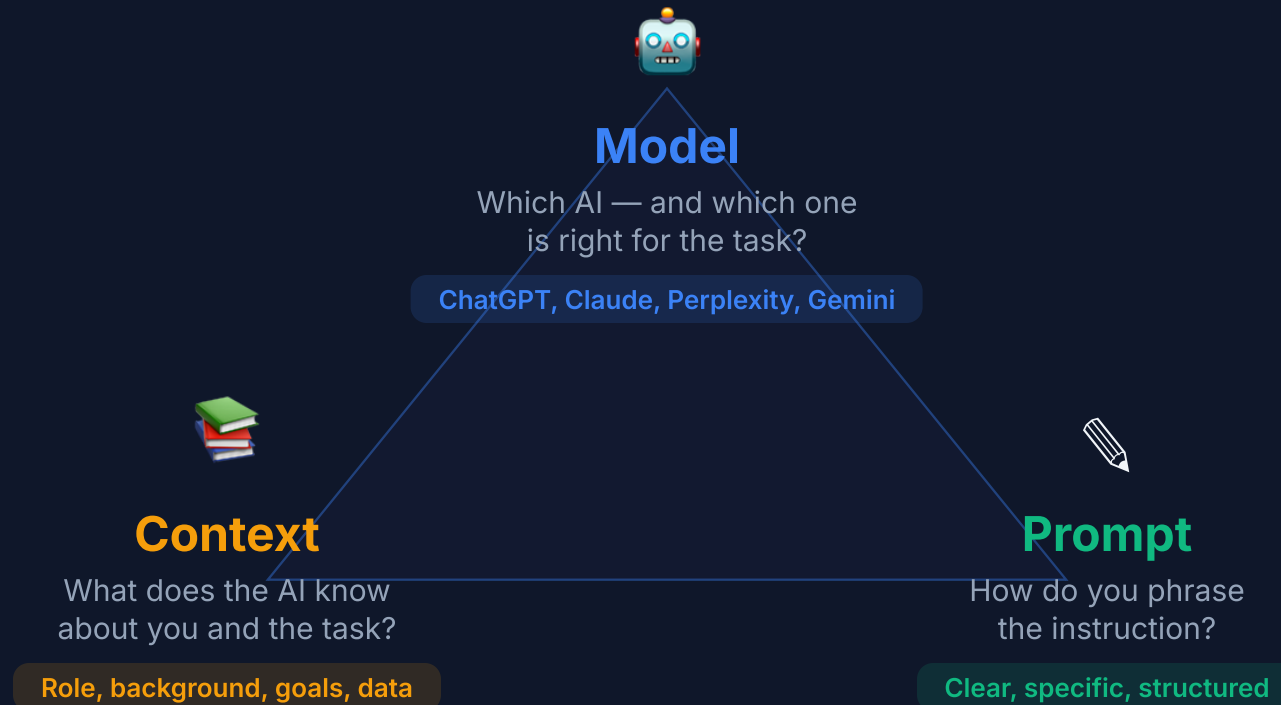
A New Kind of Tool

AI that doesn't just *talk about* solutions — it **builds them.**

This is what I use every day. And it's what I'll show you live in Block 3.

The Gen-AI Triangle

Three things determine the quality of every AI interaction.



Most people only think about the prompt. But **context** is what makes the difference.

Same task. Different prompt. Different result.

THE LAZY PROMPT

Write me an email to a customer about a delivery delay.

What the AI doesn't know:

Who is the customer?
How long is the delay?
Is there a solution?
What tone — formal or friendly? What language? What's the relationship?

Result: Generic, useless email.

THE GOOD PROMPT

You are a customer service manager at an industrial coatings company. Write a professional but warm email to a long-term client (David Miller, plant manager at ArcelorMittal). His order for STARK coating sleeves is delayed by 2 weeks due to a raw material shortage. We're offering a 5% discount on the next order as compensation. Keep it under 150 words. Tone: direct but relationship-focused.

Result: Specific, usable, ready to send.

Let AI write your prompts.

The fastest way to a great prompt is not writing it yourself. Ask Claude or ChatGPT to craft one — for any tool. Save this text and paste it whenever you need a better prompt.

PROMPT CREATOR

paste into Claude / ChatGPT

I want you to become my Prompt Creator. Your goal is to help me craft the best possible prompt for my needs. The prompt will be used by you, ChatGPT. You will follow this process:

1. First, ask me what the prompt should be about. I will give you my answer, but we will improve it through continuous iteration by going through the next steps.
2. Based on my input, you will create three sections:
 - a) **Revised Prompt** – your revised prompt. It should be clear, precise, and easy for you to understand.
 - b) **Suggestions** – details I should add to the prompt to improve it.
 - c) **Questions** – relevant questions about additional information you need to improve the prompt.
3. The prompt you provide should take the form of a request from me to be executed by ChatGPT.

Faster: a system shortcut

Set a text-expansion shortcut so the whole prompt fires from a single trigger.

Type: `promptx`
Get: → the full prompt, pasted.

macOS: System Settings → Keyboard → Text Replacements. Windows / Cross-platform: Raycast, espanso, TextExpander.

Cross-tool: Use Claude to write a prompt for Perplexity. Use ChatGPT to write a prompt for Claude. Prompts travel.

4. We will continue this iterative process: I provide additional information,

Three Modes

One platform — three ways to work.

MODE 1

Chat

Like ChatGPT — with Claude asking smart follow-ups. The familiar entry point.

MODE 2

Cowork

AI as a collaborative partner.
Connect a local folder.
Analyze documents together.

MODE 3 — RECOMMENDED

Code

AI in the terminal. Reads files, writes code, executes tasks, tests its own work — autonomously.

Data Protection — A Task, Not an Excuse

The Fear

- "If I put our coating formulas into AI, competitors could get them"

- "Our trade secrets (like STARK) could leak"

- "Customer data might end up somewhere it shouldn't"

The Reality

- These concerns are **valid** — not paranoia

- But the answer isn't "don't use AI"

- The answer is: **use it right**

It's a **task to solve**, not a reason to stop.

The Separation Principle

Three layers — your data never has to touch the AI

1

The Agent Knows Nothing

No company knowledge is stored **inside** the AI itself. The agent is a blank engine — powerful, but empty.

2

Processes Stay Local

Company processes live in **structured files** — on your machine or in a secured company cloud. The AI reads them when needed, but doesn't store them.

3

Sensitive Data via API

Financial data, customer data, trade secrets — connected via **secure interfaces** (APIs). The AI can work with results without ever seeing the raw data. *This is a target architecture — not the default. Out of the box, AI sees everything you put in. That's why: architecture first.*

You don't dump everything into the AI.
You **architect what it can access.**

Your advantage isn't just the tools

Every agency will have Claude. Every competitor will build great websites.

The tools are the baseline — not the edge.

COMMODITY

Everyone gets these

- The models
- The tools
- The workflows

ONLY STALTEC HAS THIS

Your context

- ◆ Your projects, clients, decisions
- ◆ Your tone, your standards
- ◆ Years of Staltec know-how

Your moat isn't the AI. **It's the context you feed it.**

What a company achieved with this

A German consulting firm (~15 people) built an AI OS for their finance processes.

Monthly close BEFORE

4-5 days

- Manually check budgets
- Copy-paste notifications
- Open each invoice individually
- Manually match payments
- Maintain Excel financial plan



Monthly close AFTER

< 1 hour

- Let AI skills run
- Give approvals
- Review results in dashboard
- Work on other things meanwhile
- Run scenarios by voice

~40 invoices/month. March 2026: All invoices done by the 2nd of the month.

THE SHORTEST VERSION

Get Claude. Start today.

What you saw today is a new kind of tool. **You don't learn a tool by planning it** — you learn it by using it.