

Build **infrastructure** for quality content at scale



Hi, I'm Noy 🙌



I manage this community

Building with marketing teams here

+10
Years
In Marketing

10x
Marketers

10x
Company

Quick Recap

Most marketing teams are still trying to write the perfect prompt.

But when five people on the same team prompt the AI five different ways, the result isn't scale, it's twenty outputs that look like they came from twenty different companies.

The fix isn't a better prompt, but the infrastructure behind it:

A set of building blocks: **Context, instructions, and tools.**

With the right infrastructure in place, the AI pulls from them on its own, so every output stays consistent no matter who's behind the keyboard. Here's what that looks like in practice.

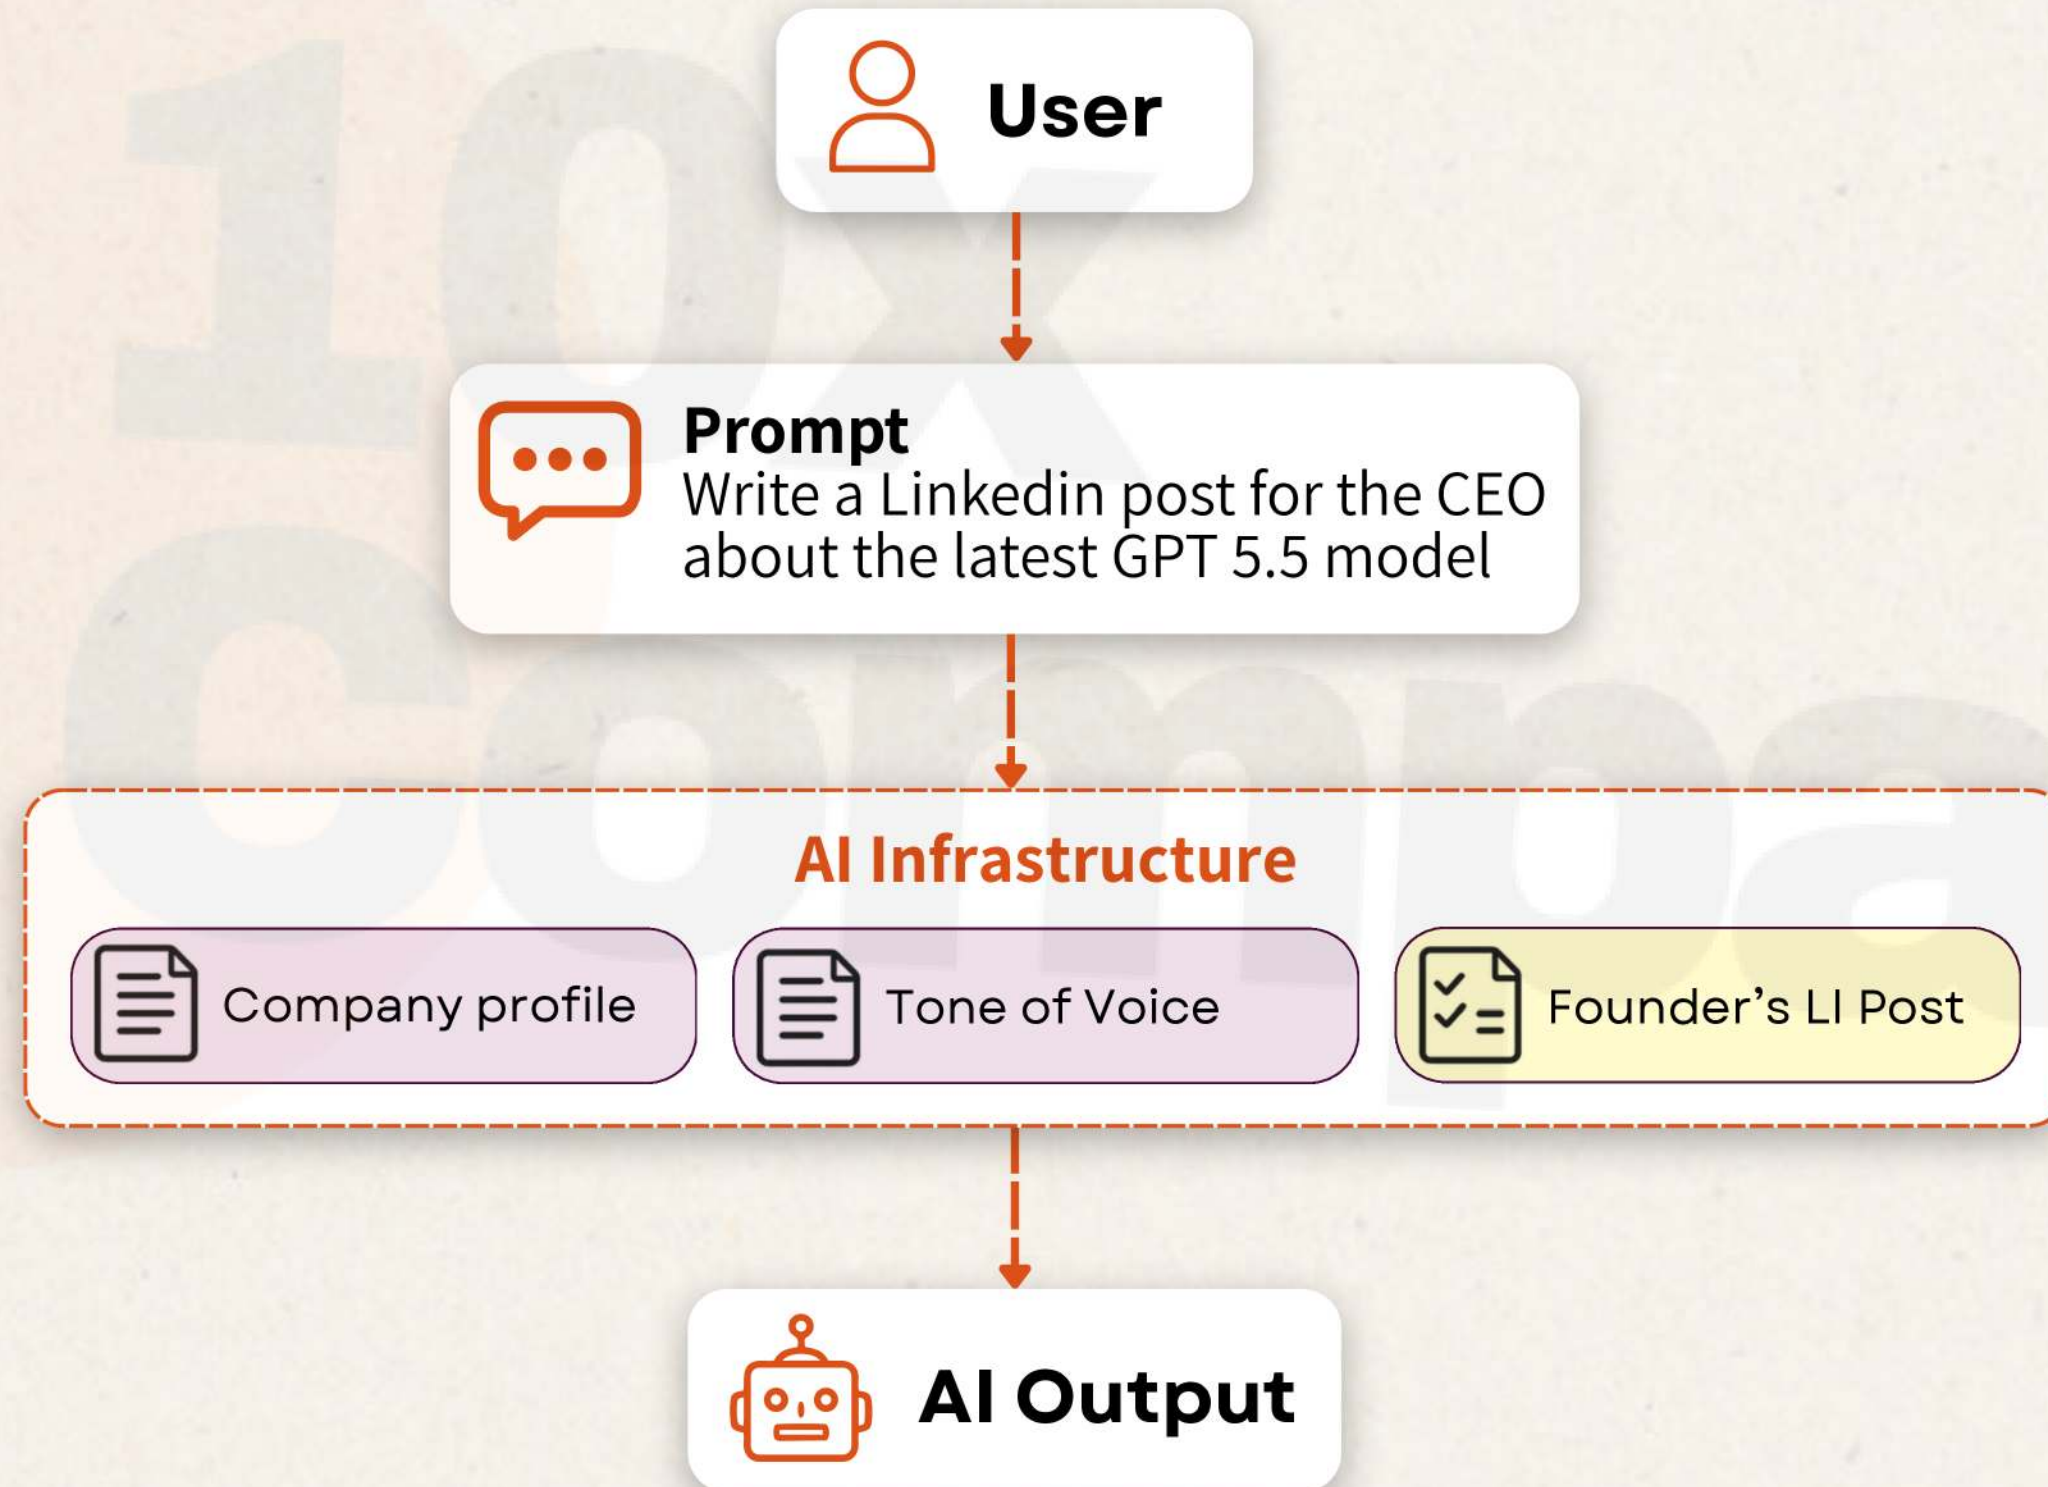
Enter, building blocks: **The infrastructure**

Tools | **With** | Gmail | Hubspot | Apify

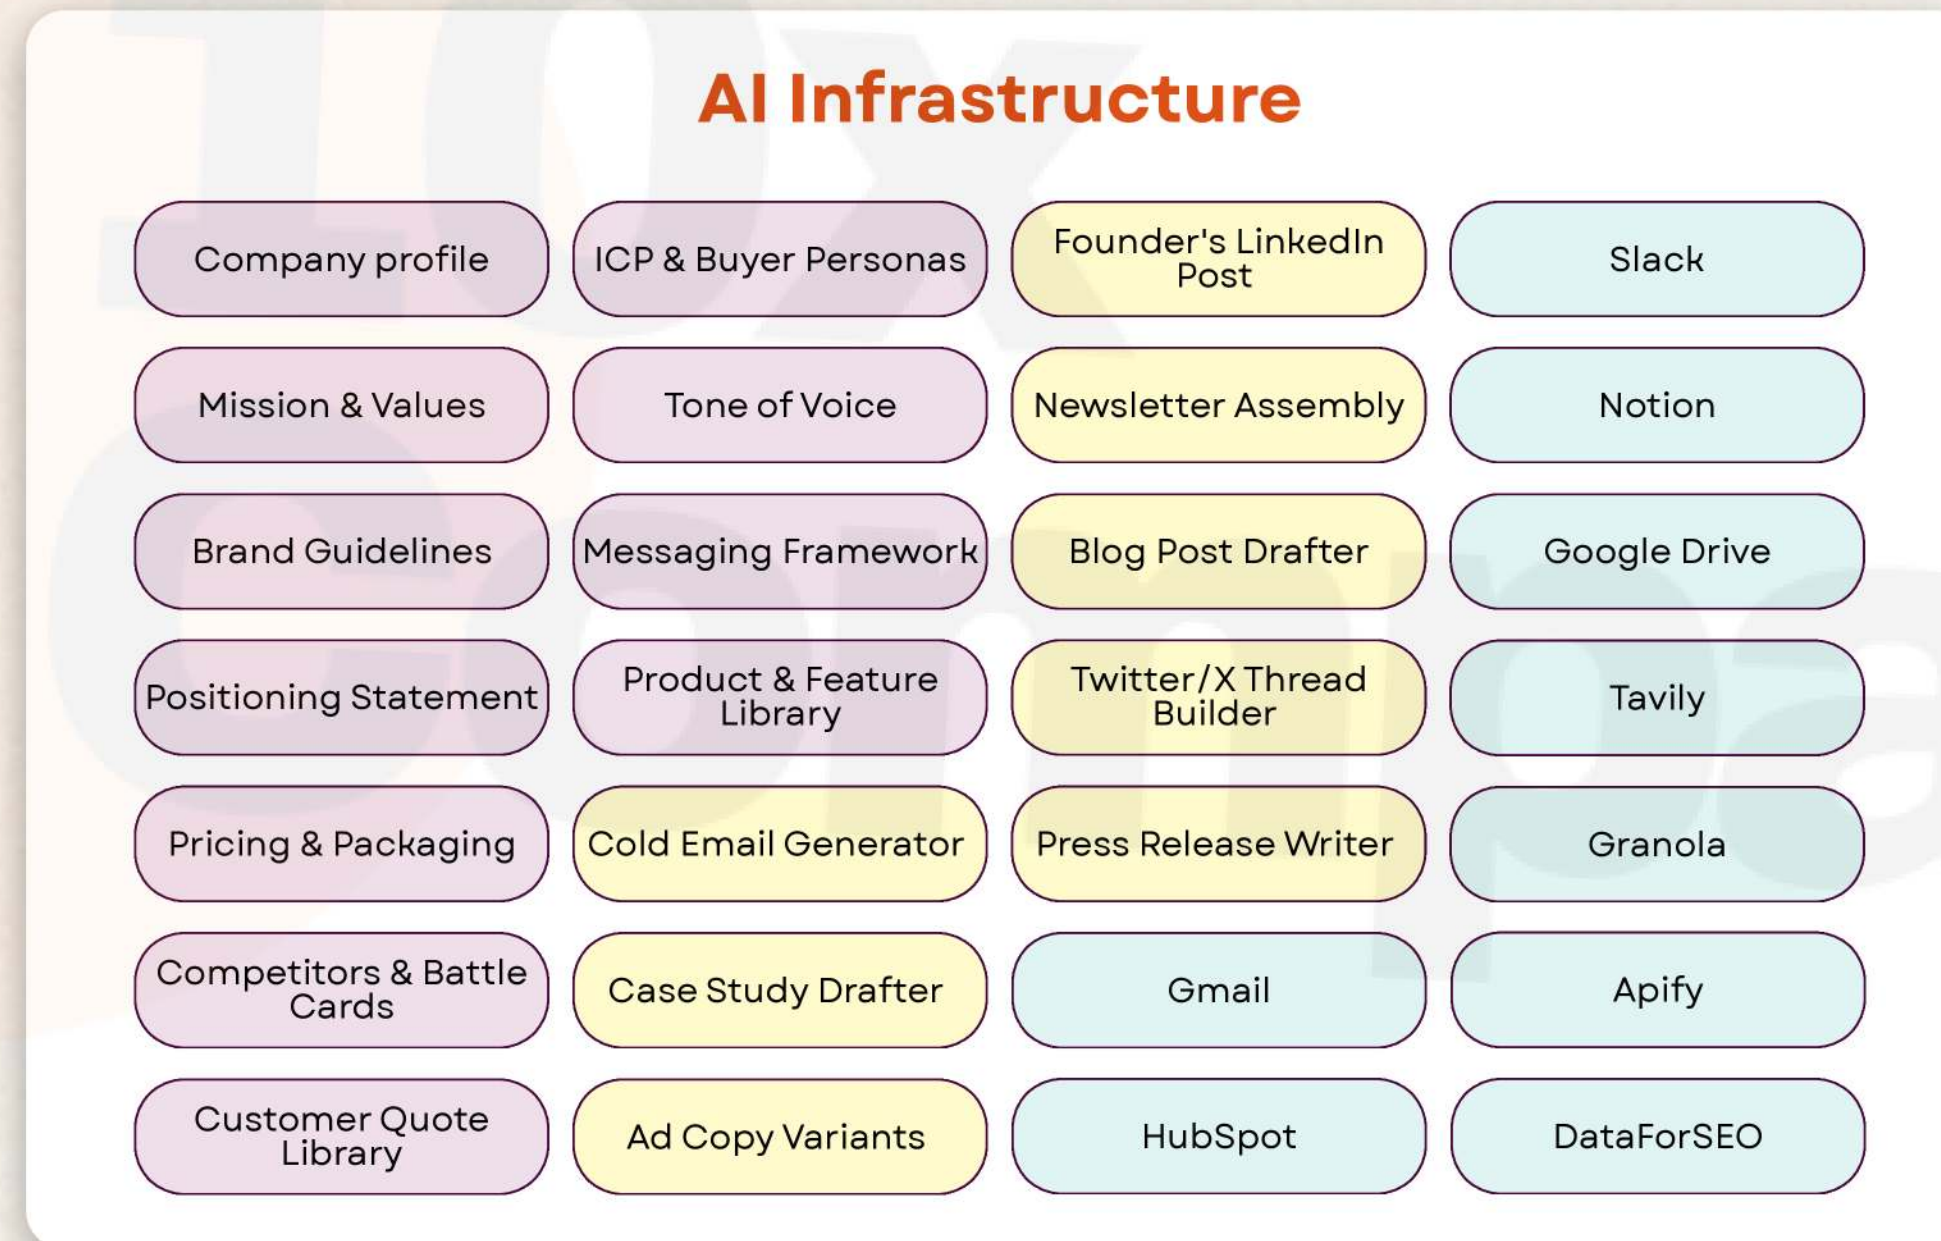
Instructions | **How** | Founder's LI Post | Weekly newsletter | Competitive Research

Context | **What** | Company profile | Competitors | Tone of Voice

Back to basics, but with **the right setup**



So how does such an **infrastructure** look like?



Example 1: Founder-led social post

Write a LinkedIn post for the CEO about the latest GPT 5.5 model

Thinking...

Two hours after GPT-5.5 launched, a CISO I respect texted me one line: "I think we just lost a year."

Every frontier model release does the same thing. It shrinks the gap between "an attack is theoretically possible" and "an attack is in the wild".

Three years ago that gap was months. Last year, weeks. With 5.5, we're going to see it measured in days. The part most people are missing is the economics. Spear-phishing personalization is about to become cheaper to run at scale than generic phishing, and that flip happens this

Reply...

+ Ask

Opus 4.6

Instructions

Founder's LinkedIn Post

Context

Company Profile

ICP & Buyer Personas

Tone of Voice

Tools

Tavily

Notion

AI Infrastructure

Company profile

ICP & Buyer Personas

Founder's LinkedIn Post

Slack

Mission & Values

Tone of Voice

Newsletter Assembly

Notion

Brand Guidelines

Messaging Framework

Blog Post Drafter

Google Drive

Positioning Statement

Product & Feature Library

Twitter/X Thread Builder

Tavily

Pricing & Packaging

Cold Email Generator

Press Release Writer

Granola

Competitors & Battle Cards

Case Study Drafter

Gmail

Apify

Customer Quote Library

Ad Copy Variants

HubSpot

DataForSEO

Example 2: The monthly newsletter

Put together this month's newsletter

Thinking...

Subject: The 4.7-minute breach is here. So is your homework.

Body:
Hi {{first_name}},
Quick one before you scroll past: median dwell time for AI-driven attacks just dropped under 5 minutes. If your detection stack is still calibrated to a 2023 baseline, you're already losing time you can't get back.

This month, three things worth your attention:

1. The numbers nobody's putting in a deck yet. Our threat lab pulled fresh data across 1,400 customer environments. AI-driven attack dwell time is now a median of 4.7 minutes. We wrote up what changed and where it's headed.

Reply...

+ Ask

Opus 4.6

Instructions

- Newsletter Assembly

Context

- Company Profile
- ICP & Buyer Personas
- Tone of Voice
- Product & Feature Library

Tools

- HubSpot
- Granola
- Notion

AI Infrastructure

Company profile	ICP & Buyer Personas	Founder's LinkedIn Post	Slack
Mission & Values	Tone of Voice	Newsletter Assembly	Notion
Brand Guidelines	Messaging Framework	Blog Post Drafter	Google Drive
Positioning Statement	Product & Feature Library	Twitter/X Thread Builder	Tavily
Pricing & Packaging	Cold Email Generator	Press Release Writer	Granola
Competitors & Battle Cards	Case Study Drafter	Gmail	Apify
Customer Quote Library	Ad Copy Variants	HubSpot	DataForSEO