VertoTalks

Making AI Work Like Part of Your Marketing Team: How We Built VertoBrain

What we learned building VertoBrain, and what any marketing team can copy

Yasen Lilov & Lachezar Dimitrov · 29.10.2026 · CampusX

VertoDigital

LET'S GET TO KNOW EACH OTHER

We'd love to hear a little about *you*

A quick show of hands:

- Who uses Claude at work every week?
- Who's on a Team or Enterprise plan?
- Who has connected Claude to their own data?



48 people joined us. Were you one of them?

VERTOTALKS #18 · JANUARY 2026

We're Betting on AI: Inside Our AI & Automation Strategy

JAN 29

Thursday 29 January · 10:00–11:30

Campus X, Sofia

WHAT WE COVERED

The strategic bet – why we're investing in AI now

Our vision – where AI fits into B2B marketing

The 24-month roadmap – internal → packaged → revenue

Lessons learned – what's working and what isn't

WHAT WE SAID IN JANUARY

From “AI is an initiative” to “AI is how we work”

01

A dedicated team

Dedicated leadership, full-time people and a tooling budget.

02

Scale what we build

From R&D and pilots to tools people use every week.

03

Bring everyone along

Adoption across every team, not only the enthusiasts.

And one lesson we put on screen: build vs buy is a decision you keep making.

WHY MOST AI EFFORTS STALL

Almost everyone uses AI. Very few have scaled it.

88%

of companies use AI somewhere in the business

7%

have scaled it across the organisation

WHERE IT BREAKS



Sources: McKinsey, The State of AI · BCG, Where's the Value in AI?

THE LENS FOR TODAY

Process → Context → People

01

Process

Pick the right battles

Write down how the work is done, then automate what repeats.

Today: skills and routines that carry the team's method

02

Context

Your data is the difference

Data and know-how are what make a general model useful to you.

Today: VertoBrain, one door to live client data

03

People

Blocker and multiplier

Tools don't adopt themselves. People who use them every day do the convincing.

Today: how 40+ people use it day to day

What it taught us

A strong model with half the context gives you confident, wrong answers.

THE STORY

Someone on the team asks Claude: *“Which campaigns are at risk of overspending this month?”*

The answer comes back from live data, for that one client only.



A model is only as good as the harness around it.

WHAT A HARNESS IS

Everything you'd give a smart new hire on day one

A new hire gets their logins

01 · THE FACTS

Access to the real data

…the handbook

02 · THE KNOW-HOW

What good looks like, written down

…the team's jargon

03 · SHARED LANGUAGE

Your words mean one thing

…the checklists

04 · REPEATABLE WORK

The same steps, every time

…limits on what they can touch

05 · GUARDRAILS

What it can see, touch and say

…and someone keeping it all current

06 · UPKEEP

The handbook never goes stale



We built ours big.

You can build yours small.

PART 1 OF 6 · THE FACTS

One connector in front of all our client data.

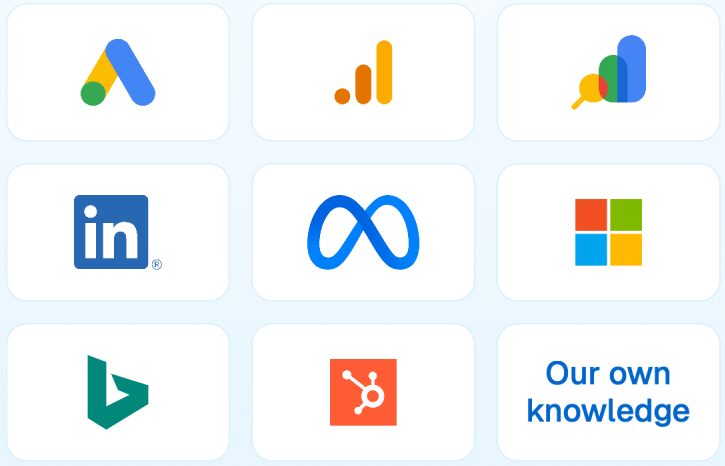

Ads, analytics, search, CRM and our own knowledge

ONE DOOR

VertoBrain

Confirms the client first, then routes each request to the right source. Read-only by default.

Claude Team

Chat, Cowork, Claude Code and routines all use the same door.

Also behind the door: OpenAI Ads, DataForSEO, ScreamingFrog and more.

ONE REQUEST, START TO FINISH

What a normal request looks like

THE ASK

"How is paid search going for [client] this month, and which campaigns are at risk?"

01 VertoBrain confirms which client

02 Live data comes back for that client only

03 Claude answers, showing the numbers it used

At risk, most severe first

Campaign	Spend	Conv	CPA (Sept)	Issue
EMEA brand + non-brand	$6.1k	1	$6,069 ($575)	Conversions collapsed from 10 to 1. Rank-lost IS is 58% and CTR fell to 3.3% from 4.2%. Budget isn't the constraint.
[redacted] demo	$19.3k	14	$1,382 ($892)	CPA up 55%. Impression share is 13%, with 50% lost to budget and 38% to rank.
[redacted] demo experiment	$3.6k	3	$1,198	77% of impression share lost to budget. The experiment is competing with its parent.
[redacted] brand general	$23.2k	20.9	$1,108 ($901)	33% of total spend. Brand IS is only 29%, with 35% lost to budget and 36% to rank. CTR fell to 5.7% from 7.6%. It is running at 1.5× its daily budget, so brand terms are being taken by competitors.
[redacted] demand gen TOFU	$4.8k	0	n/a	Running at 1.9× its $90 daily budget. Zero conversions.
EMEA expansion	$2.0k	0	n/a	Zero conversions, and Sept had 1. Low dollars at stake.

THE FACTS · THE BREAKDOWN

Give Claude the facts, not your memory of them

WHAT WE HAVE

- One door to 16 sources
- Live data, trimmed to what matters

WHY IT MATTERS

- Exports go stale and get cut
- No data? It guesses, confidently

DO IT YOURSELF

- Connect every tool, even unlisted ones
- Learn what each one sends back

Next: both, step by step

DO IT YOURSELF · THE FACTS

Tool not in the directory? Add it yourself.

Most tools now publish their own connector.

ALWAYS VERIFY THE SOURCE

Only the vendor’s own URL

01 Search “[tool] MCP server”

02 Copy the vendor’s URL, e.g. mcp.tool.com/mcp

03 An Owner adds it in Organization settings

04 Everyone signs in with their own login

DO IT YOURSELF · THE FACTS

Interview your data before you rely on it

You can’t change what a connector returns. You can teach Claude to read it.

01 Ask something you already know

02 Ask which tools and fields it used

03 Compare with the platform. Explain every gap

04 Save a data card per source

Write down what good looks like

WHAT WE HAVE

- Skills for our core work
- Built from flows people already ran

WHY IT MATTERS

- Same method, whoever asks
- Juniors start from the senior’s checklist

DO IT YOURSELF

- Give every SOP step its reason
- Package them as shared skills

Next: what that looks like

DO IT YOURSELF · THE KNOW-HOW

Write the *why* next to every step

STEP ONLY

“Exclude brand terms from the non-brand report.”

Ask why, and Claude invents a reason.

STEP + WHY

“…because brand terms inflate CTR and hide how non-brand is really doing.”

Juniors learn the why. Claude spots the case that doesn’t fit.

Claude carries your seniors’ experience to everyone.

Say it the way you'd say it to a colleague

SOMEONE SAYS	WITHOUT SHARED LANGUAGE	WITH OUR SKILLS AND CONNECTORS
"Build a topic cluster map for [client]."	A generic mind map, different every time.	Our topic cluster map, our way.
"How is [client] doing on paid search?"	A web search for three similar firms.	That client's live data, straight away.

Teach Claude your team's words

WHAT WE HAVE

- Client names resolve to their accounts
- Skills named after deliverables: audit, brief, dashboard, QBR

WHY IT MATTERS

- People talk to Claude like a colleague
- No "which one did you mean?"

DO IT YOURSELF

- Write down clients and deliverables in your words
- Test with real requests

Next: the three pieces

DO IT YOURSELF · SHARED LANGUAGE

Make your team’s words trigger the right work

01

A client register

Names, nicknames, IDs, lookalikes to ignore

02

Skills in your words

Claude picks a skill by its description

03

Test with real requests

Ten messages from your team chat

PART 4 OF 6 · REPEATABLE WORK

Our own platform: ai.vertodigital.com

01 Fill in a short form to start the audit

02 A routine runs for 15–50 minutes in its own environment, every step visible

03 Out comes a near-finished audit: findings and a deck to download

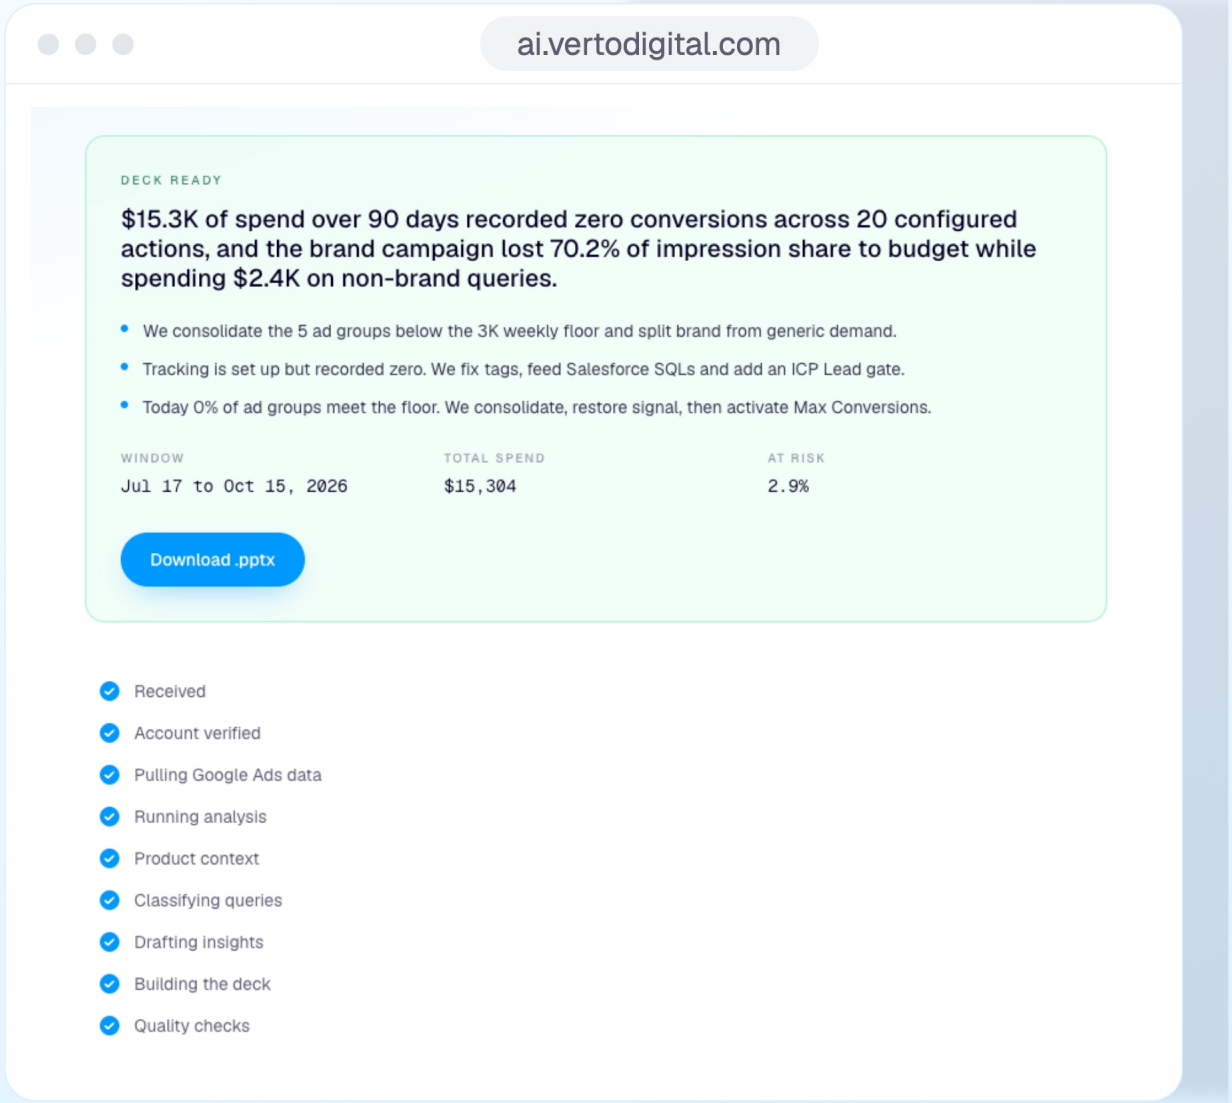

Turn work you repeat into work that runs

WHAT WE HAVE

- Audit routines, started from a form
- Own environment, built-in checks

WHY IT MATTERS

- Same steps and checks every run
- Nobody babysits a 30-minute job

DO IT YOURSELF

- Turn your best session into a skill
- Schedule it once it’s stable

Next: what to tell Claude

DO IT YOURSELF · REPEATABLE WORK

Your best session is a skill waiting to happen

Long session that finally works? Ask Claude to turn it into a skill. Include:

01 The goal and output format

02 Inputs: client, period, sources

03 Steps and decisions

04 Edge cases you hit

05 Checks before delivering

06 Then test it in a fresh chat

Scheduled task: from your desktop app · Routine: in the cloud

Decide what Claude can see, touch and say

WHAT WE HAVE

- Client first, every request
- Read-only; personal data removed

WHY IT MATTERS

- Client data never crosses over
- The team can trust the output

DO IT YOURSELF

- Limit access at the source
- Make Claude check and show its work

Next: five ways to do that

DO IT YOURSELF · GUARDRAILS

Make Claude prove its work before you see it

01

Read-only at the source

A view-only login for every tool

02

Rules it must follow

“Missing data? Say so.
Never estimate.”

03

“Before you deliver”

Every skill ends with checks to report

04

A second agent checks

A validator in the plugin reviews it

05 · WHAT WE DO

Every deliverable comes with its working

Calculations, sources and assumptions, in a separate sheet

PART 6 OF 6 · UPKEEP

Keep the handbook current

WHAT WE HAVE

- A weekly agent keeps the docs current
- 1–2 hours of upkeep a week

WHY IT MATTERS

- Stale instructions, confident wrong answers
- Only as good as its last update

DO IT YOURSELF

- Live connectors over uploaded files
- A named owner for every skill

Sometimes Claude alone isn’t enough.

V/PROBER · A PILOT

What does ChatGPT search for before it answers your buyers?

01 Add your buyers' questions

02 ChatGPT answers each many times; we record every search

03 Searches become a must-cover checklist

WHY IT LIVES OUTSIDE CLAUDE

Other models · many runs · fixed steps · cost control

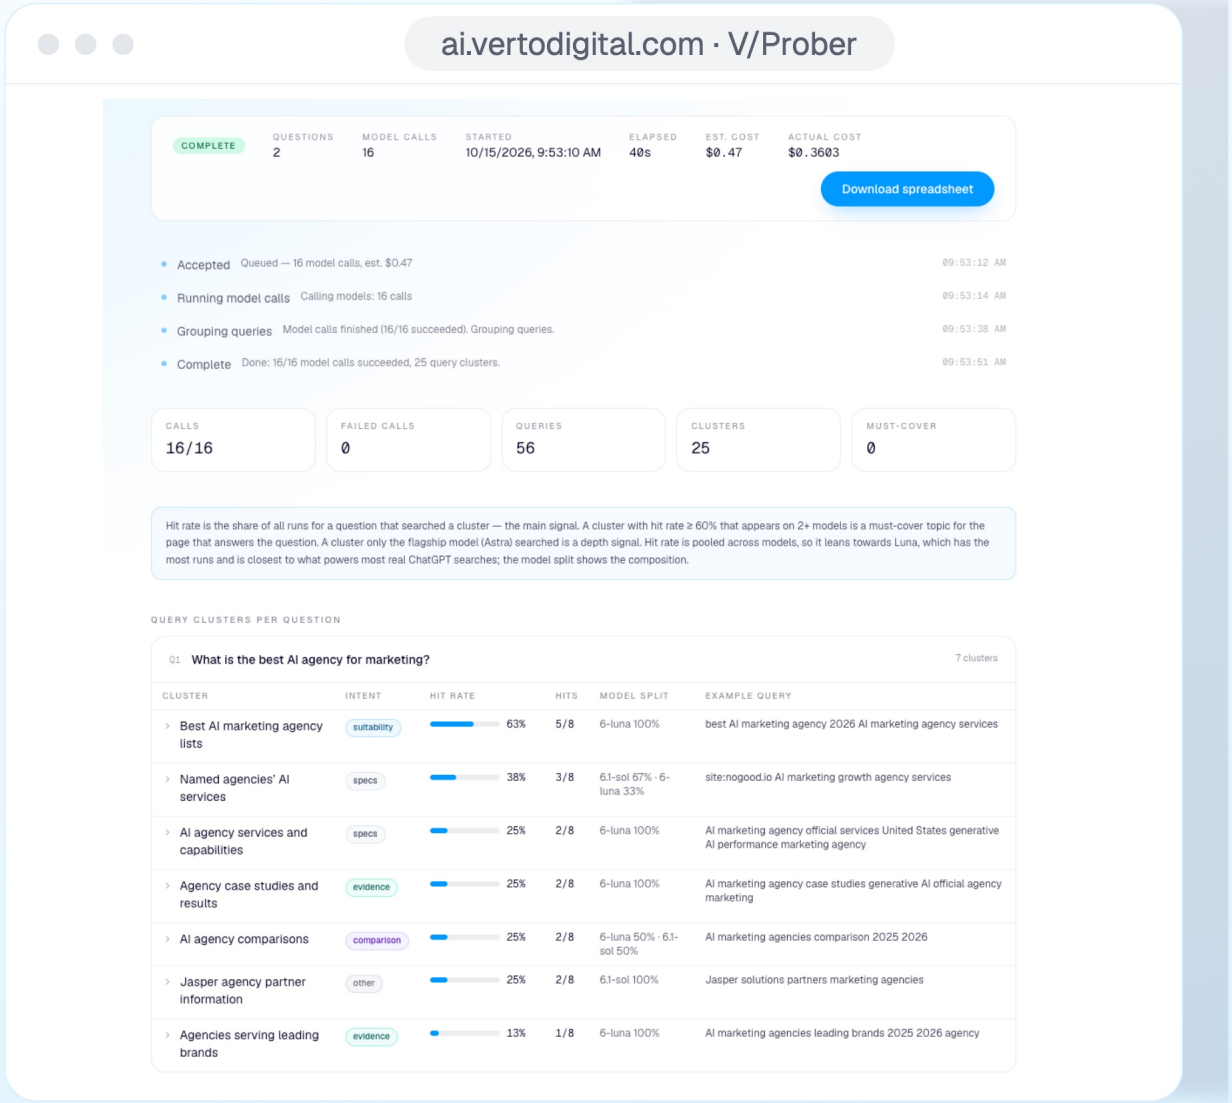

So, did it work?

40+

PEOPLE USE IT EVERY DAY

16

DATA SOURCES BEHIND ONE DOOR

2

PEOPLE BUILT IT, IN ABOUT TWO MONTHS

15–50 min

FROM REQUEST TO NEAR-FINISHED AUDIT

vs weeks by hand

1–2 h

OF UPKEEP A WEEK



1

DOOR TO ALL OF IT

THE LESSON WE'D PASS ON

A good harness is built from real work, not from a library

80%+

of the skills and plugins we started with are hardly used at all

HOW WE STARTED

Skills from online libraries and generic ideas

WHAT PEOPLE ACTUALLY USE

Built from weekly tasks, with the people who do them

TAKE IT HOME

Your harness checklist

The facts	Connect your tools · a data card per source
The know-how	SOPs with the why · shared skills
Shared language	A client register · skills in your words
Repeatable work	Best session → skill → schedule
Guardrails	Read-only · checks before delivery · show the working
Upkeep	Live over static · an owner for everything

PRACTICAL TIPS

Six habits that turn good answers into great ones

01

Make it interview you

“Ask me what you need to know first.”

02

Name the reader

Who reads it? What will they decide?

03

Get a cold second opinion

Fresh chat: “Review this as a sceptical client.”

04

Hand over, don’t scroll

Long chat? Handover note, fresh start.

05

Ask what caused the mistake

Then fix the instruction, not the answer.

06

Explain it to a newcomer

Find your gaps. Know what kind of task it is.

VertoTalks

THANKS, OVER TO YOU

Questions & *stories*

Ask us anything, or tell us how your own team uses Claude.

vertodigital.com Yasen Lilov & Lachezar Dimitrov

VertoDigital

Appendix

Backup slides for technical questions

Five layers between a question and client data

01 **One entrance.** Client first, every time.

02 **Our own connectors.** Built and vetted by us.

03 **Look, don’t touch.** Read-only by default.

04 **No personal data.** Stripped automatically.

05 **Platform protections.** Claude Team + Cloudflare.

Same steps, every run.
Same checks, every run.
Zero usage for whoever asks.

Built-in checks · runs on the creator’s plan

A 60-SECOND REFRESHER

Three Claude features this story runs on

01

Connectors

Claude reads your tools

Live data instead of pasted exports

02

Skills

Your know-how, packaged

Everyone does it the team’s way

03

Routines

Work that runs on its own

On a schedule or a trigger

HOW WE BUILT IT

Two people, about two months

Our knowledge

Claude learned our clients and methods from our own documents

Google Ads

The first platform, because that's where most of the team works

The rest

A few more platforms, then all of them

One door

Every platform moved behind a single connector

Launch

Live for the whole team, about a week later

APPENDIX · ONE WORKAROUND WORTH STEALING

Making Cowork follow our steps instead of improvising its own

USUALLY

A fresh script every time, a little different each run

OUR WAY

Our tested scripts, handed over in the session

01 Someone asks for a task

02 The matching skill switches on

03 It asks VertoBrain for the scripts it needs

04 Our internal services connector fetches them

05 Data and scripts land in the Cowork session

06 Cowork does the job exactly as we designed it

The model still does the thinking. The steps that must never vary are ours.

Claude Team covers most of it. We fill three gaps.

01

Other models

Claude Team is Anthropic-only

Some tasks suit other models

02

n8n flows

A front end Claude doesn’t have

Forms, status and run history

03

Fixed steps

Chat can’t promise them

Prototyped in n8n. Also our lab.

APPENDIX

How the pieces fit together

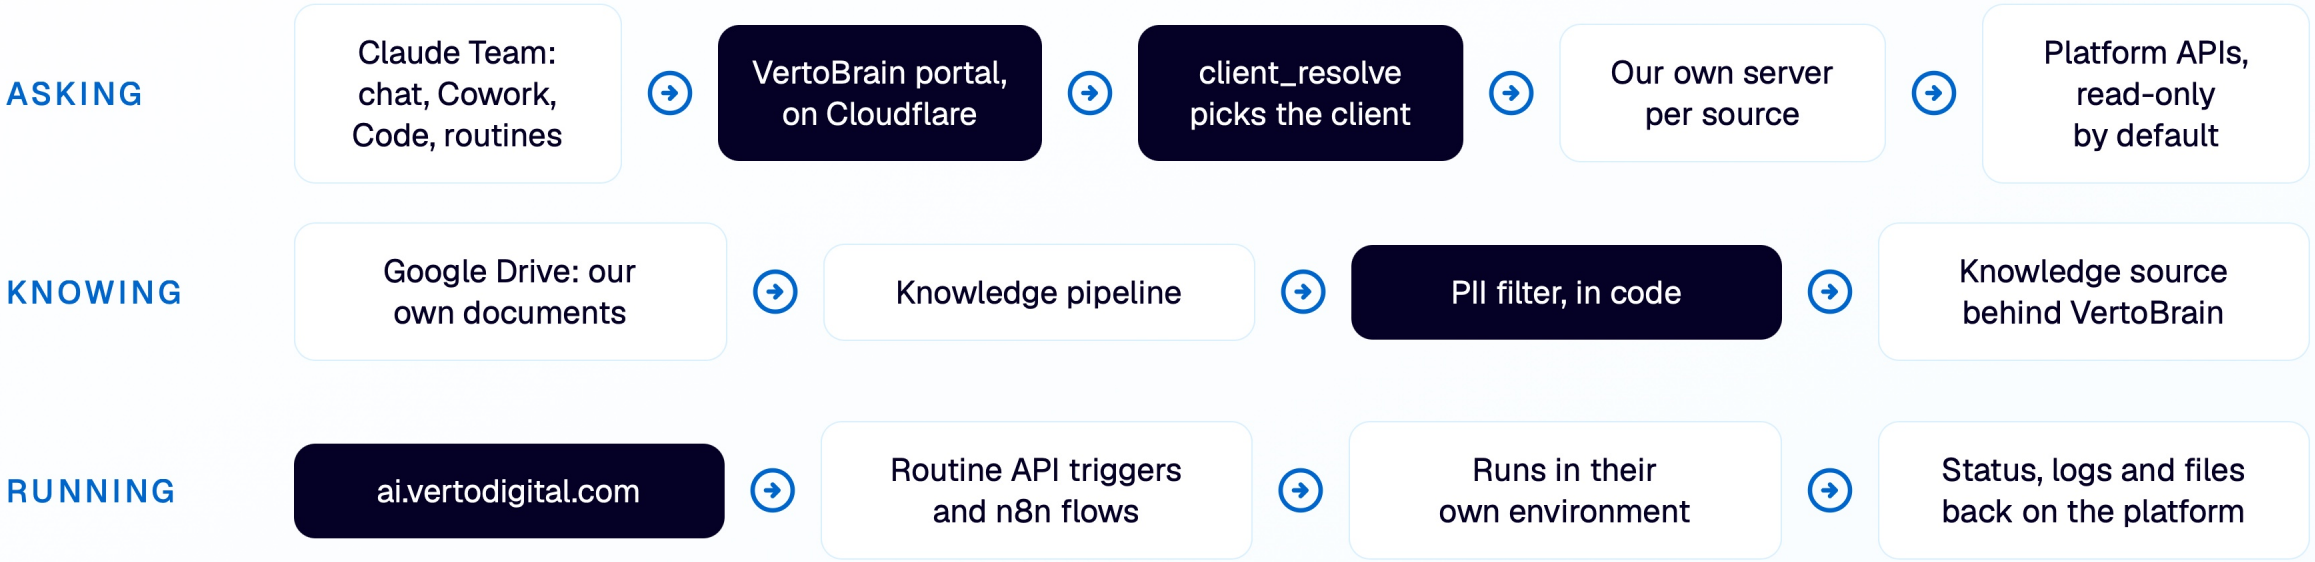

PLACEHOLDER Check every box against the live setup, and add anything missing (e.g. other model providers).

APPENDIX

Why one portal: the stateless MCP update

WHAT CHANGED

Requests no longer need a session with one server, so one portal can route them all.

BEFORE

One connector per platform, each added and authorised on its own.

AFTER

One connector, one sign-in, every source behind it.

Easier adoption
One connector to add and one sign-in, for everyone

Better cross-source work
One conversation pulls ads, analytics and CRM together

Stronger security
Sources talk over internal bindings, with no publicly exposed URLs

Less room for client data to slip
Only two tools are available at the door

Less context used
Only those two tools load by default, not every source's tools