




Accelerate agent adoption across every team, **with every agent you approve.**

Cartesian is agentic adoption and acceleration for your agent ecosystem. It sits inside Claude, ChatGPT and Copilot, matches the task someone is working on to an approved agent, and offers it right there, from the day the agent is approved.

Every team uses your approved agents from the day they ship, nobody rebuilds one that already exists, and you see the return as it happens.



PRIVATE Runs inside your own cloud account, and reads the conversation there to make the match. **SECURE** Independently audited to SOC 2. **DATA** Your agent code, your prompts, and your agent activity stay in your environment. **ACCESS** Your identity provider decides who can call what.



Your agent ecosystem, organised around your people and their work.

Cartesian does the organising. It packages your approved agents and puts them to use wherever your people work.

OFFERED WHILE THEY WORK

Cartesian matches the task someone is working on to an approved agent and offers it inside the tool they're using, whether or not they knew the agent existed.

EVERY APPROVED AGENT, IN ONE PLACE

Your own agents and the vetted third-party ones, each bundled with the agents it depends on and published once.


RUN AND LOGGED

The agents run in order inside your own cloud account, and every recommendation and every run is logged to the team that used it.



Your tools build, govern and index an agent, **but that's where they stop.**

YOUR TOOLS TODAY




CARTESIAN STARTS HERE



Vendor positions as published April to August 2026. Cartesian reads a registry. It does not replace one.

Approved agents, the moment someone needs them.



Illustrative scenario, not a customer result.

Cartesian 2026

THE CHALLENGE

The team that tested your agents already uses them. But how do you get the rest of the organisation using them?

HOW IT WORKS

Cartesian reads the conversation in Claude, ChatGPT or Copilot, works out the task, and finds the approved agent in your catalogue that does it. Nobody needs to know the agent's name.

RUN THE AGENTS

They run it or pass, and both are logged to their team. The agents run in order, and two years of behavioural research decide which agent gets offered and when.

ALSO AVAILABLE IN ChatGPT · Microsoft Copilot · Gemini Enterprise ·
Notion · Cursor · and more to come

Adoption that keeps up as agents, models and teams change.

AGENTS CHANGE

New versions ship. Old ones get retired. What was approved last month isn't always what's running today.

MODELS CHANGE

Swap the model and what the agent can do changes with it. The catalogue has to know that happened.

YOUR WORK CHANGES

A new standard, a new tool, a team that just got reorganised. The match has to keep up.

PEOPLE JOIN

Someone starts this week with no idea what exists, and nobody reads an announcement from six months ago.

Announcements and training reach whoever's in the room, on the day. Cartesian offers the agent every time the work comes up, whenever that is.

The return when your agents are in use **from day one.**



Adoption at speed

New agents in use across every team the day they're approved.


People get the agent while they're doing the work, so nobody waits for the next training session or goes looking for the guide.



Inventory

Every approved agent, where it is, who can call it.

When your regulator or your board asks what AI you are running, you answer from the catalogue.



Reuse

One published agent serves every team.

Teams stop rebuilding an agent that already exists, so there are fewer builds and fewer fixes.



Measure impact

Every offer, every run, and every gap.

You see what got used, what actually solved the task, and what people asked for that no agent covers.

Illustrative scenario.

Agents rely on agents. Package the whole thing.

Underneath, Cartesian is a package manager, for the people who build and publish the agents.

ALL AGENTS IN A WORKFLOW CHAIN


Example agent chain: a report agent calls a data agent for the numbers and a QA agent to check them, and they have to run in that order.

PACKAGED ONCE

Package the three agents as one, declare the dependencies, and publish it to your catalogue. Every team runs the approved chain from the client tools they work in.

AUGMENT WITH EXTERNAL AGENTS

Third-party agents you have approved sit in the same catalogue, so nobody wastes time building what an existing approved agent already does.



Illustrative scenario. Nothing here is a customer result.



See adoption as it happens, and what it's worth.

See usage against real demand, across every team.
Cartesian logs every agent run the moment it happens.

WHAT'S LOGGED


Every step of every run, including the ones that failed: tokens, tool calls, and the model it escalated to.

WHAT COUNTS AS SOLVED

It isn't the agent saying it finished, it's an external check that passes, or the user who accepts the result and does not come back.

WHAT YOU COMPARE IT TO

The baseline you agree before anything is configured, measured again over the same window later.



Key reporting data includes cost per solved task, usage gaps, churn reports, and user-led agent requests. Illustrative figures.

Three milestones, one team, **no fee.**

START

MILESTONE 1

MILESTONE 2

MILESTONE 3



Agree the baseline

You see: the baseline sheet and the measures you agree to prove.

Publish and run

You see: your first agent found and run from Claude, and 5 to 10 agents in the catalogue.

Measure reuse

You see: the reuse count and engineering hours against the baseline.

Review results

You see: cost per solved task against the baseline. Keep it running at the price agreed up front, or switch it off.

CARTESIAN GIVES

Your private catalogue and an Outpost in your own cloud account, yours from the first milestone and staying there. We surface your agents in the tools we agree on, and you get an engineer plus a founder every fortnight. You get a say in the roadmap, and we agree the price and the reference terms before you start.

YOU GIVE

One owner and four hours a month. One team, and 5 to 10 real agents in the catalogue at the first milestone. Fortnightly feedback and the run data. A yes or a no at milestone 3.

NEXT STEP Pick one agent. Agree what it has to prove. Your team, your baseline, your owner, and a decision inside two weeks.

Proposed test plan, pending sign-off from both founders. The milestones are scoped to what you want to prove, not to a fixed calendar. Book the 45-minute session at cartesian.io/design-partnership.