

Your next user does not have EYES

What the agent saw on your checkout page

- ☒ "Add to basket"
a button, and it says so
- ☐ A clock icon
read as a countdown timer
- ☐ An arrow with no label
read as decoration, it was the next step

Where the research came from

We went through what Google, Anthropic, Microsoft, LangChain, Asana and Notion have shipped in their agent products, alongside the studies published in 2025 and 2026.

Twenty patterns came out of it. **Five are not about the agent you build.**

They are about whether someone else's agent can use the product you already have.

Matt Biilmann of Netlify named this **agent experience** in January 2025: the experience an agent has as a user of your product.

1: Access

Can an agent working for your customer get in at all?

Sign-in built for a phone

Blocks your customer's agent as effectively as a bot.

Bot blocking

Cannot tell an agent working for a customer from a scraper.

Who is asking	
Agent	Booking assistant
Acting for	Priya Raman
Allowed to	Read bookings, make one
Until	18:00 today
Every action lands under both names.	

What it needs

Its own identity, for a named person, at a set scope, with an expiry.

Not a shared login

Or nothing can be traced to anyone.

An agent acting for your customer **is still your customer.**

2: Context

Does it know what your product is and how it works?

Plain text documentation

Not a video walkthrough.

Copy-pasteable examples

The part it will actually reuse.

yoursite.com/lms.txt

- ✓ What this product does, in four lines
- ✓ The five things people use it for
- ✓ Links to the docs that matter
- ✓ Where the API lives

One file. No maintenance beyond the occasional edit.

An lms.txt index

One file pointing at what matters.

Structure over styling

An agent cannot work out meaning from layout.

Semantic structure carries the meaning, **or nothing does.**

3: Tools

Is there a way in that is not a screen?

An API or a command line

Something that does not require clicking.

An MCP server

The common way agents connect to tools now.

Ways in

- ☒ REST API
documented, with examples
- ☒ MCP server
tools an agent can call directly
- ☐ Skills
how those tools combine for a real job

Skills

How your capabilities combine for a job people do repeatedly.

Why they matter

Without them agents invent commands that do not exist.

In one published test, agents without skill files produced over **a thousand errors** by inventing syntax.

4: Orchestration

Can it be set off by an event, and can it hand the job on?

Events in

Webhooks, not only a person typing.

State out

A way to pass the job along with its context.

In, through, on



context travels with the job, or the handover is worthless

Agent to agent

Over 150 organisations backed a shared protocol by April 2026.

Check who is calling

Signed cards, so you know the agent is who it claims to be.

Four layers, four jobs: tools, agent to agent, your own front end, and your interface running inside someone else's app.

5: Then go and test it

Run your top three flows with an agent driving, and watch where it stops.

Same discipline as usability testing

Different user.

Three flows is enough

Sign-up, search, checkout.

Where it stopped

- 1 Sign-up needed a code from a phone
- 2 Filter only appeared on hover
- 3 Delivery date picker read as a timer

Three findings. That is the backlog.

Watch for

Icons with no label, anything needing a hover, anything needing a human eye.

The output

A list you can hand to a developer on Monday.

Synthetic testing found agents reading a clock icon beside a time field as a **countdown timer**.

Four pillars, one test

- 1 Access.** Can it get in, with its own identity and a named person behind it?
- 2 Context.** Does it know what you do and how you work?
- 3 Tools.** Is there a way in that is not a screen?
- 4 Orchestration.** Can it be set off by an event and hand the job on?
- 5 Then test it.** Three flows, agent driving, watch where it stops.

Human users still need a good interface. They are no longer the only ones arriving.

Have you ever watched an agent try to use your product?

Yes or no in the comments.

If you change one thing

Publish an llms.txt

Why this one

An afternoon of work, almost nothing to maintain, and the cheapest way to find out whether anything is already reaching you through an agent.

Then start measuring agent traffic.
You cannot design for a user you cannot see.

Go deeper

The whole pattern library

On the site

Seven families, forty-one patterns.
What each one prevents, and the
product or the paper it came from.
Free, and no sign-up.

thinkbigleaders.com/resources/agent-patterns

Or talk it through

Curious, and nowhere near a decision?
Good. Have a use case in mind? Also
good. Already building one, or want
agents inside your own product? Bring
what you have.



Thirty minutes with Think Big Leaders.

Scan the code to pick a time.