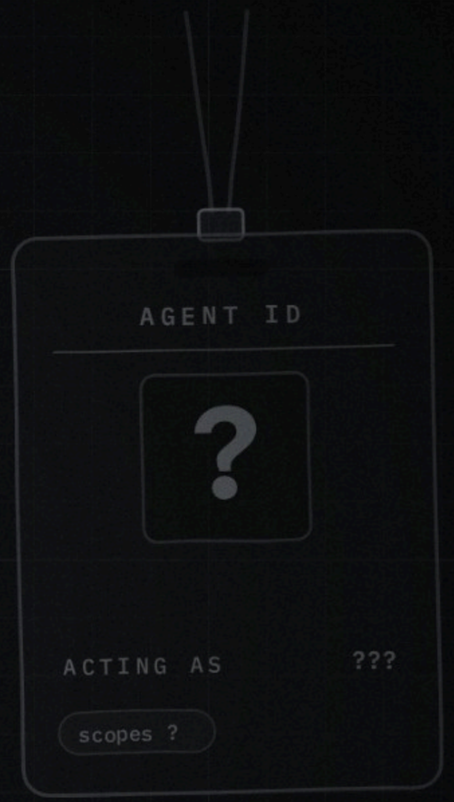
nori agentic

Agent identity is a really hard problem

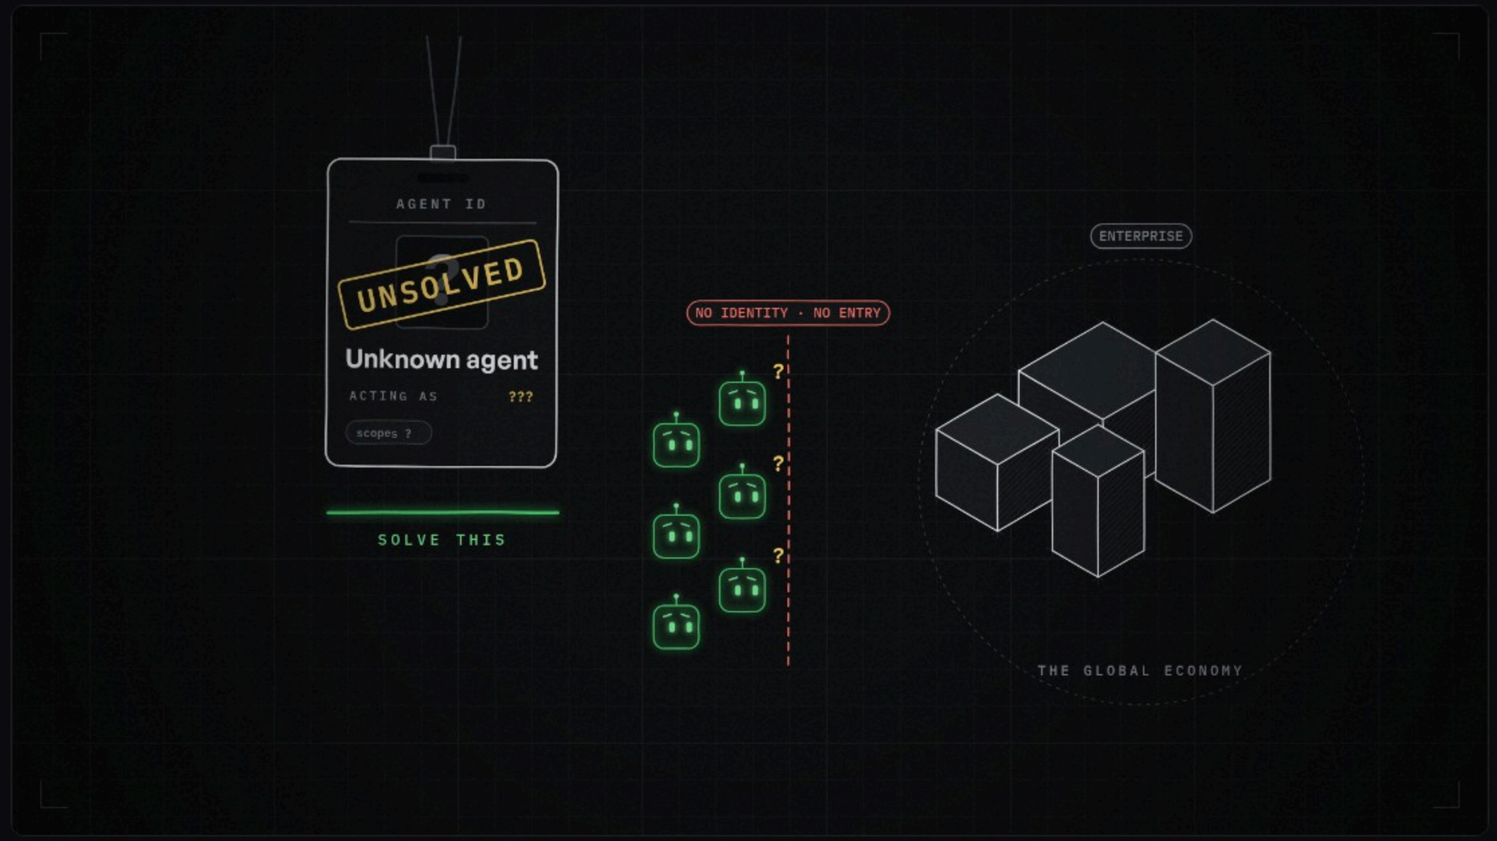


BY AMOL KAPOOR

If you want to make a bajillion dollars, solve agent identity.

As far as I know, no one has really solved this. It is a massive problem that is preventing the diffusion of AI agents across enterprise and the global economy.

src · Agent identity is a really hard problem

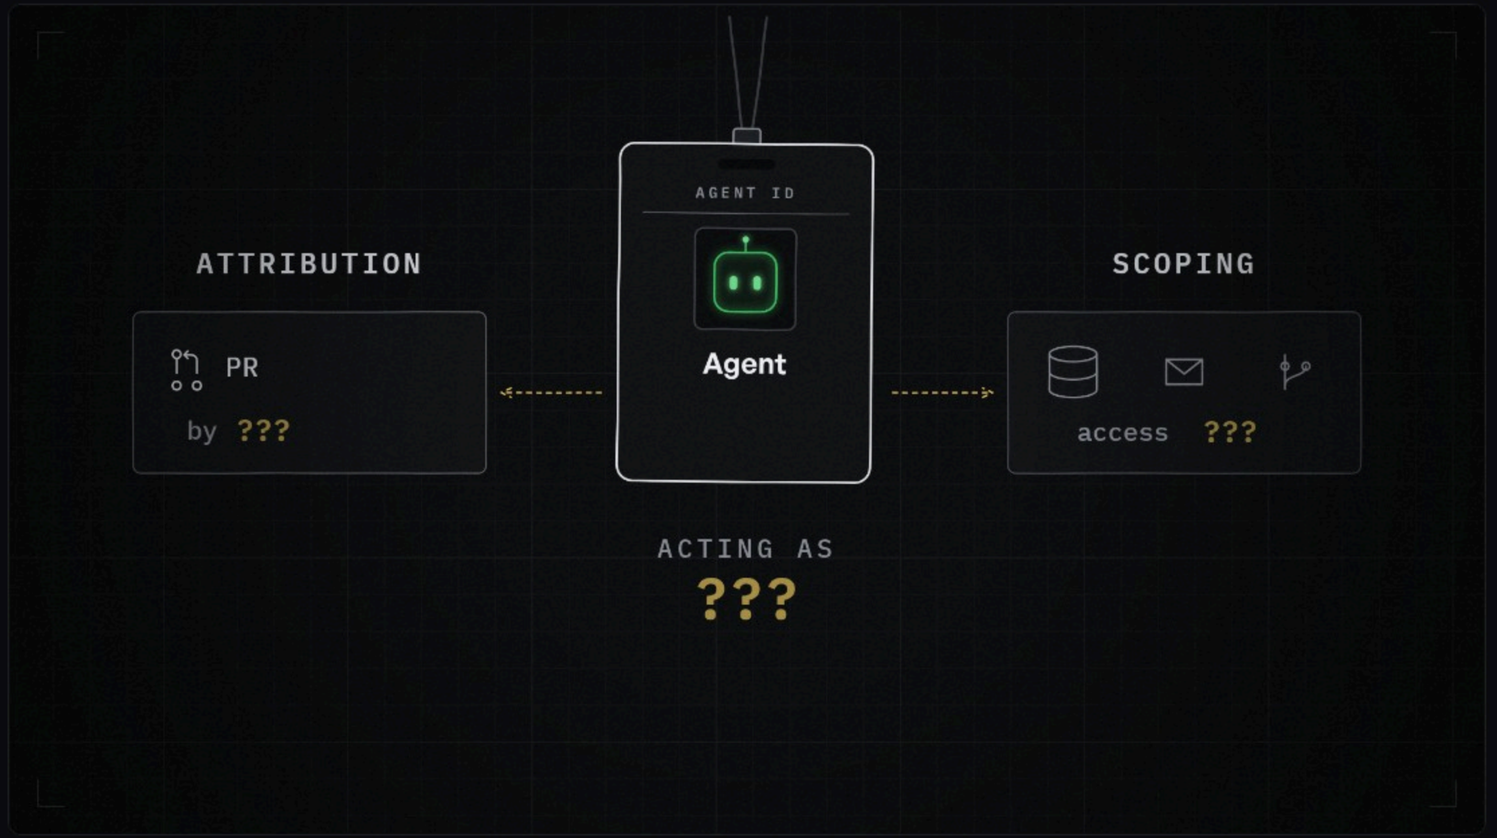


“who is this agent acting as?”

Agent identity answers what should be a simple question: “who is this agent acting as?” You can think of this as a combination of attribution and scoping.

If an agent opens a PR, who is the PR under? If an agent is given a task, what resources does it have access to? If an agent massively screws something up — hallucinates a clause in a contract or leaks a bunch of health data or takes over the world — who gets sued/fired/held responsible?

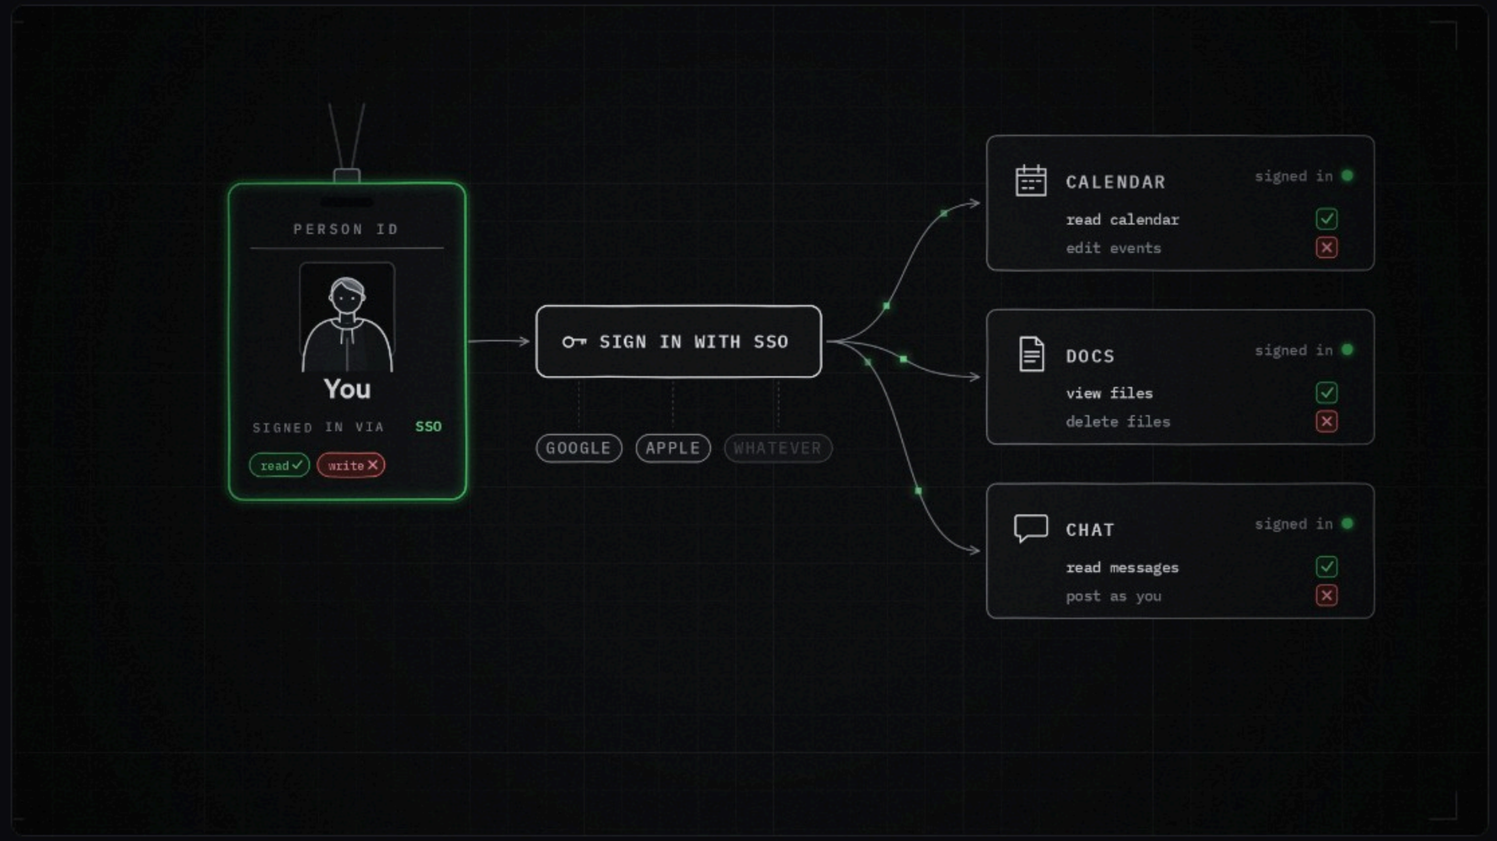
src · Agent identity is a really hard problem



We only just kinda figured out identity for people.

We can mostly seamlessly have individual people sign into a bunch of third party services using single sign on through Google or Apple or whatever. And you can generally be pretty specific about what access you're providing.

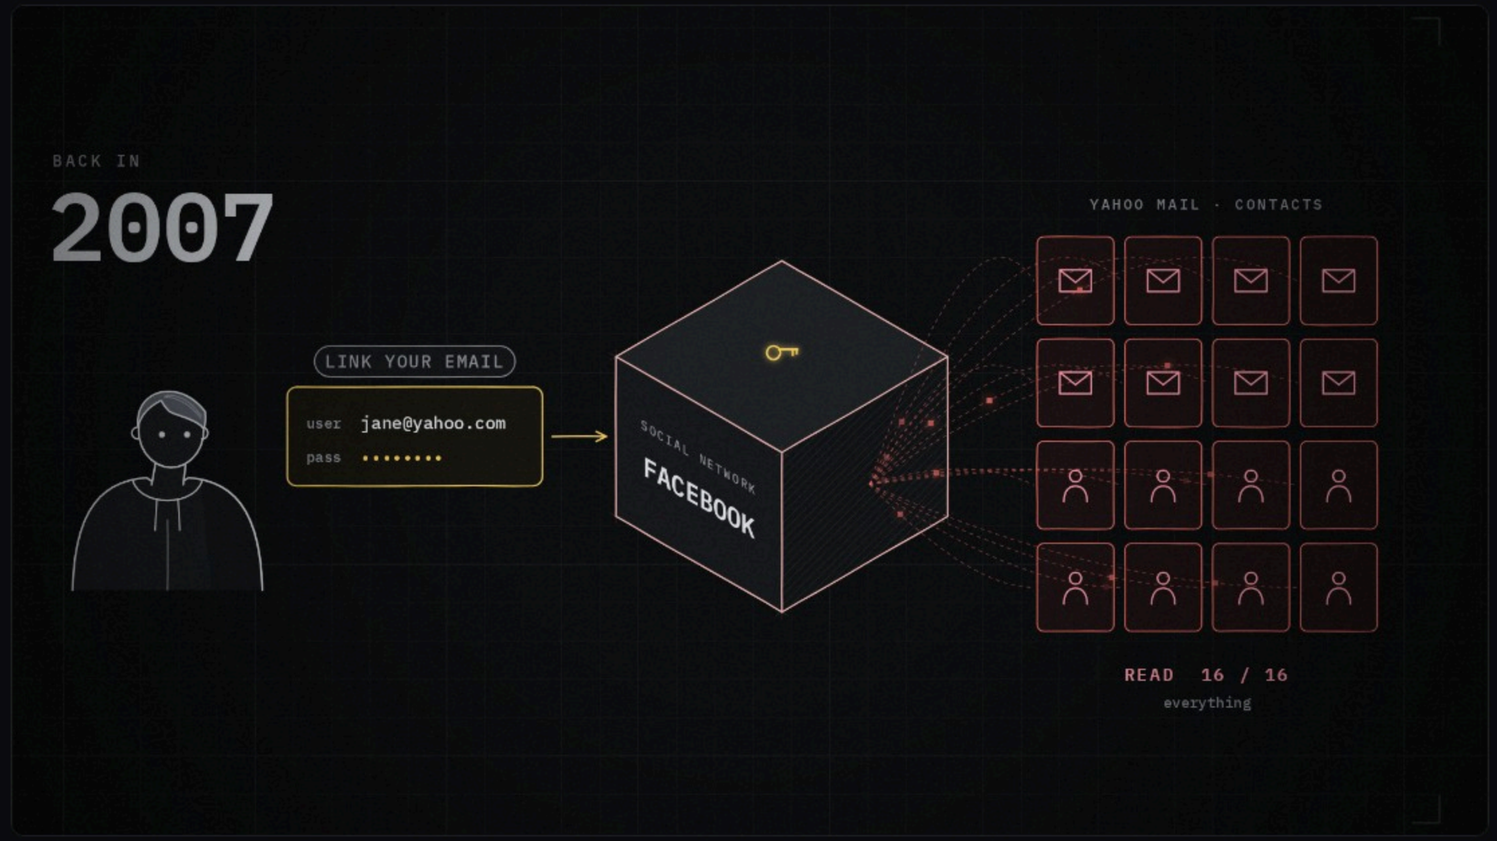
src · Agent identity is a really hard problem



SSO and OAuth were driven by bad behavior

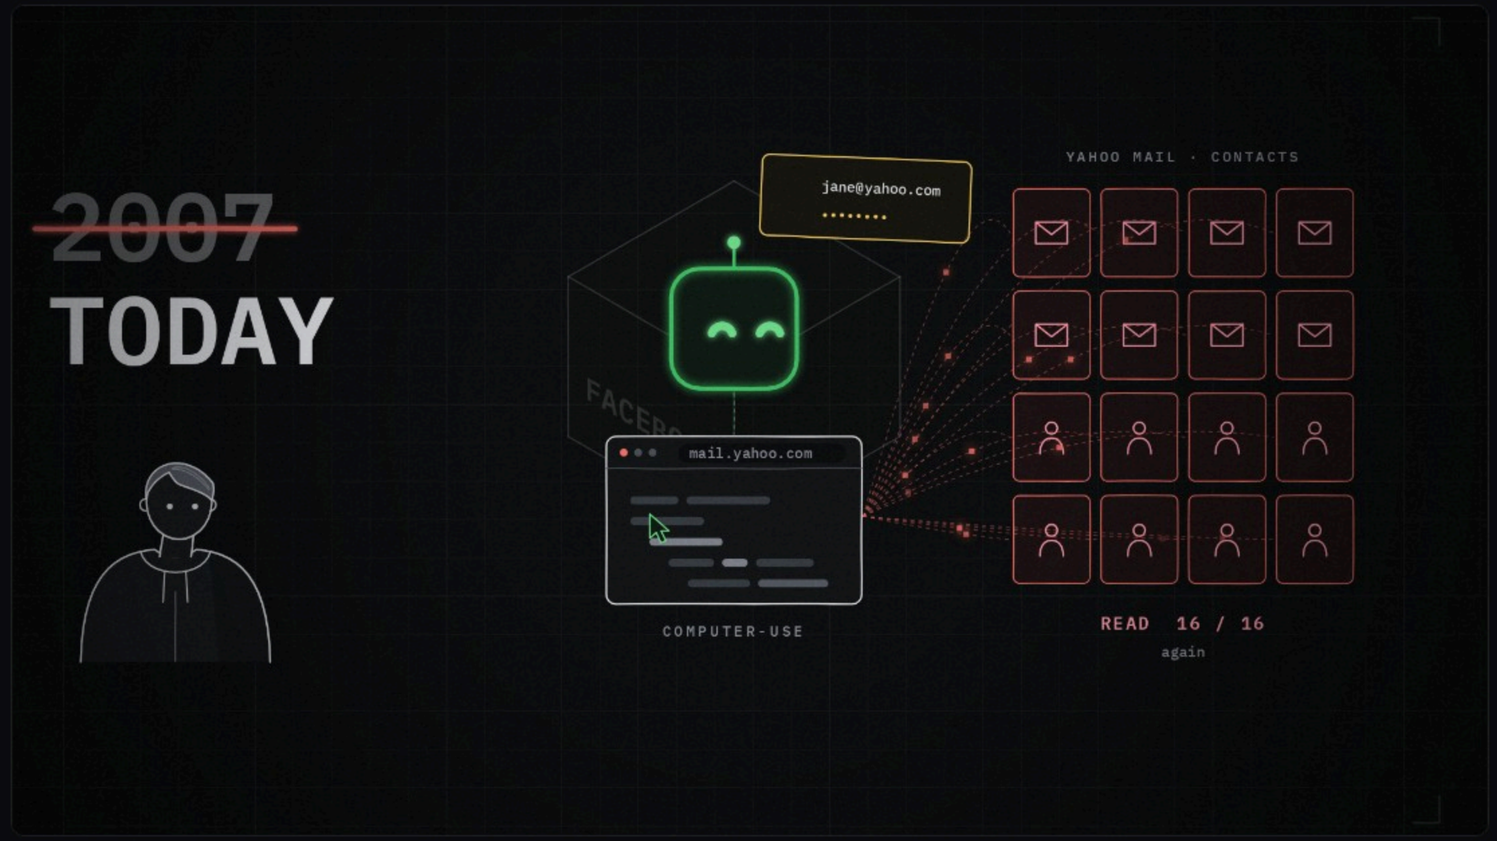
Funny enough, SSO and OAuth were driven by bad behavior from companies like Meta. Back in 2007 or so, Facebook would encourage users to link their emails and their Yahoo accounts and so on. They would do this by literally asking the user to hand over their username and password, and then would use that access to crawl literally everything they could get their hands on.

src · Agent identity is a really hard problem



Today, people use computer-use agents where they...give the agents their username and password, so that the agent can crawl through other accounts and see what they can get their hands on!

src · Agent identity is a really hard problem



We obviously need something better for agent identity. But this is a tough problem, because it’s not really clear how that identity should work.

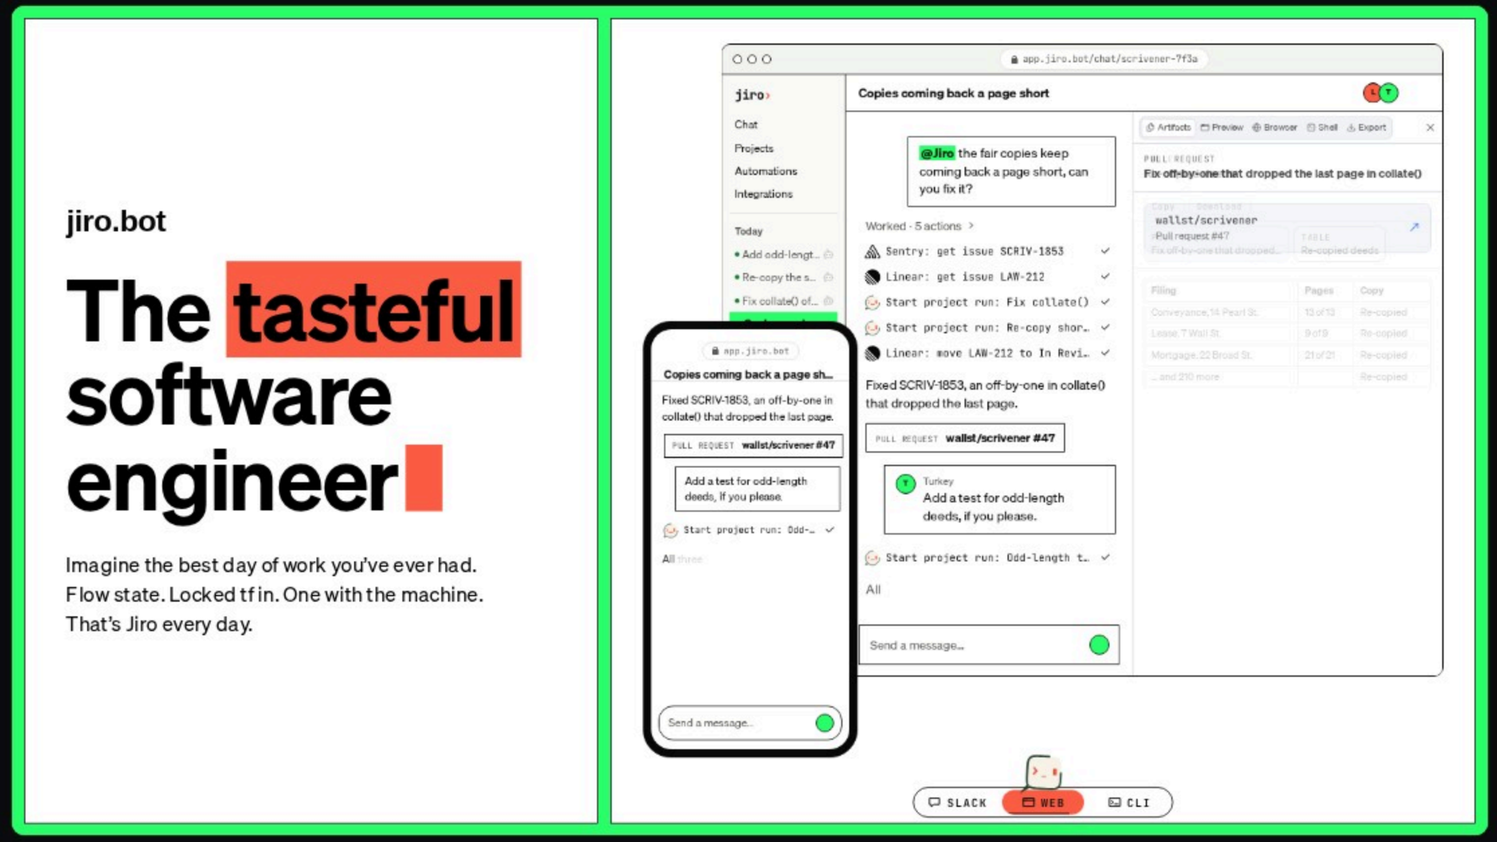
Hi, my name is Amol.

Amol Kapoor · CEO, Nori Agentic

Nori builds jiro.bot.

> A **proactive**, **tasteful** software engineer with all of our **embedded knowledge** on how to do good work.

We have worked with dozens of teams on AI infrastructure and spent a lot of time thinking about identity with our customers.



So far, we've seen three variations of how people try to solve agent identity, though they all have problems.

src · Agent identity is a really hard problem

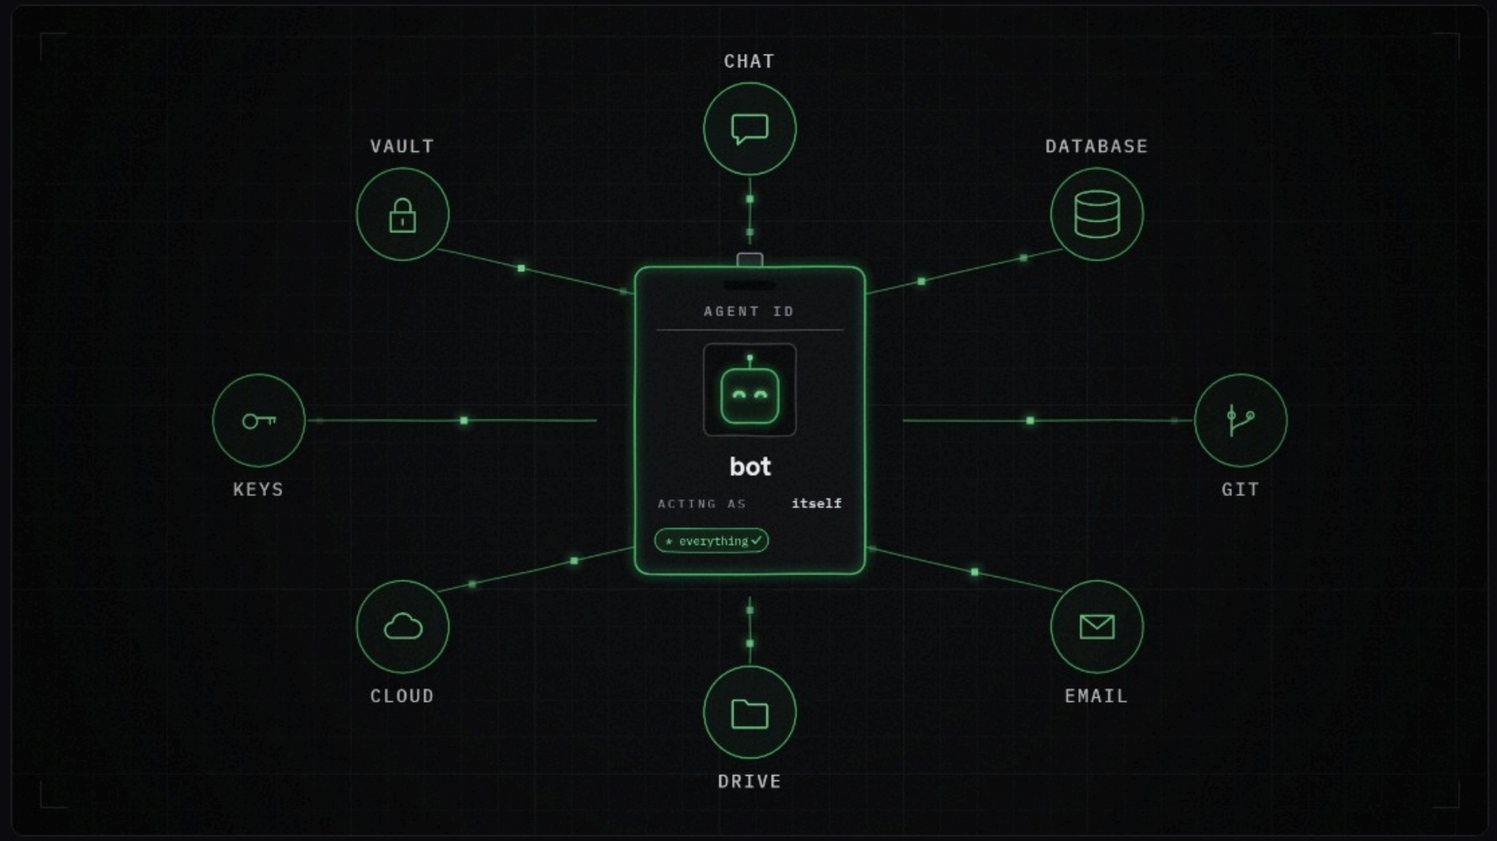


Have a single agent identity that just has access to everything

aka “we’ll do it live.” This approach is basically to ignore the problem entirely and just give the agents access to everything. Speed is the most important thing, so let’s go fast.

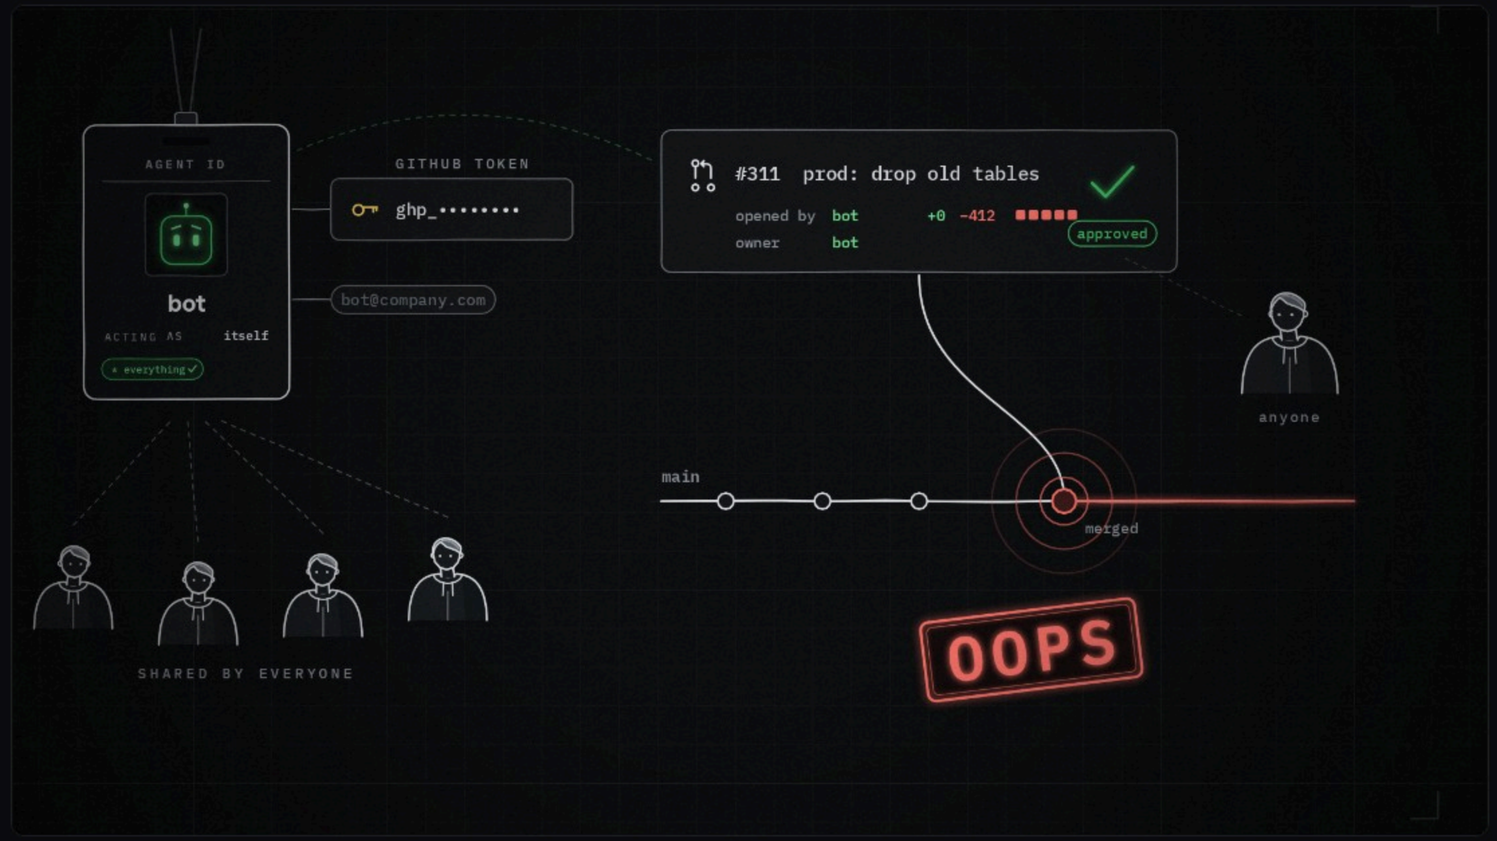
This is most popular with startups and small teams that formed in the last two years. People who don’t have significant existing revenue dependent on certain security attestations like “random people on my team cannot access sensitive data.”

src · Agent identity is a really hard problem



Generally the agent is given its own identity for tools that need it, eg its own github access tokens and its own email address. The upshot is that this makes sharing very very easy. The downside is that anyone who has access to the agent has access to anything. And there's other weird edges -- for e.g. any PR would be opened and owned by the bot identity, so anyone else can ‘review’ the PR and merge it. Oops.

src · Agent identity is a really hard problem

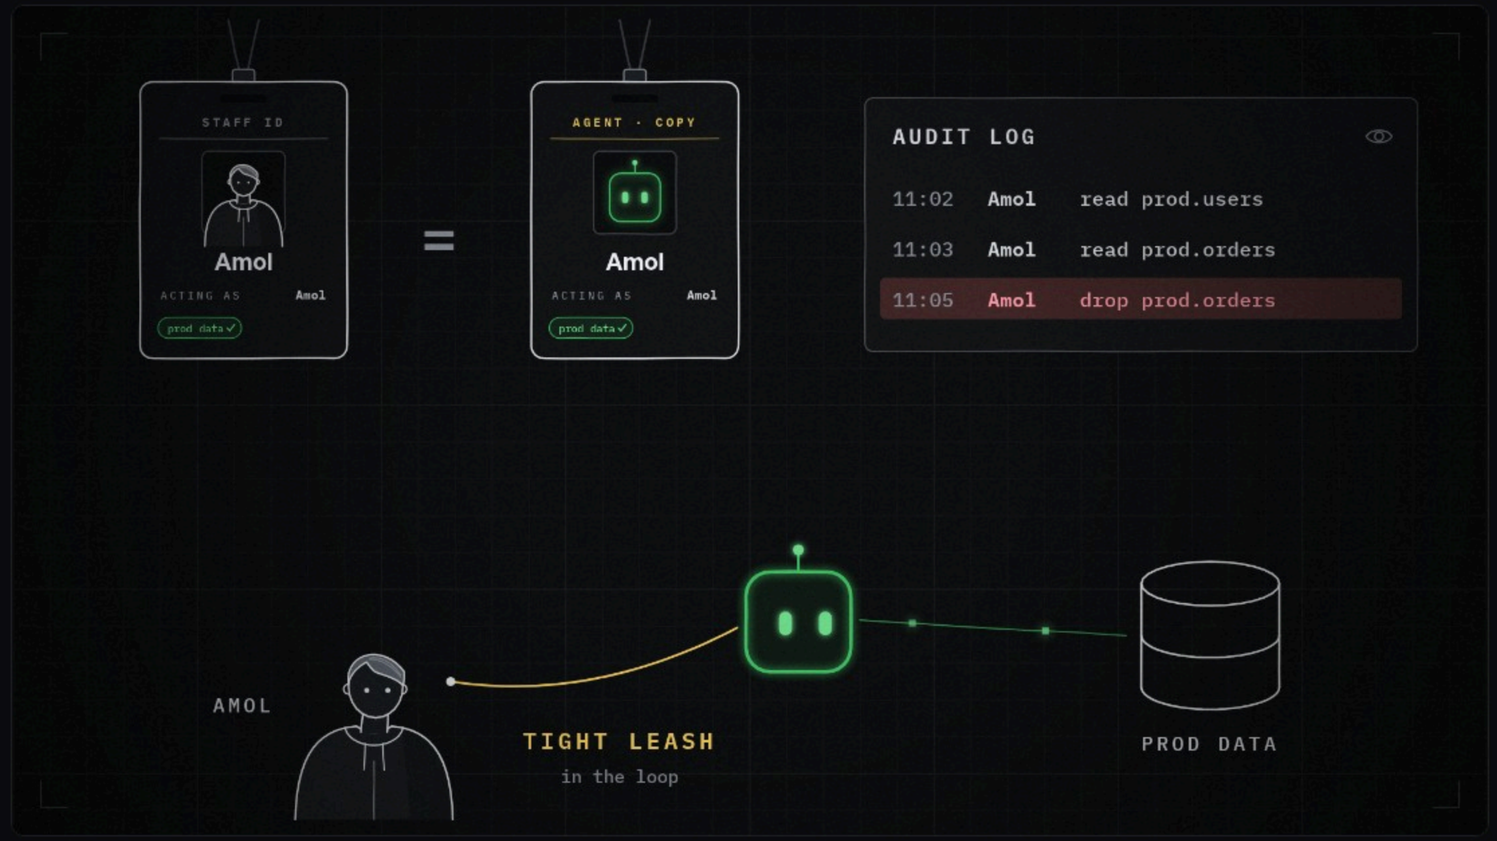


Per person identity inheritance.

In this model, the agent acts as the person who started the thread. In some sense this is great because it's really easy to reason about. If I, Amol, have production data access, then the agent has production data access. It also makes audit and accountability pretty easy. If the agent does something wrong, I have done something wrong.

But that is also the problem. This model of the world works great if you use the agent like a tool, where the human is aggressively in the loop and keeping the agent on a tight leash.

src · Agent identity is a really hard problem



This does not work at all if you are trying to have cloud agents that run on cron jobs / webhooks or even just trying to run a lot of agents in parallel more autonomously. If I can get in trouble for an agent doing something stupid, I'm not going let the agent go do whatever!

Also, like, sometimes I may not want an agent to inherit all of my scopes! I might be doing something that's reasonably dangerous, where I want the model to be in read-only mode even though I personally have more generous permissions.

Also also, the ergonomics of this are painful — you need to get everyone's individual creds into the system somehow.

src · Agent identity is a really hard problem



A bunch of agents each with a unique identity.

In this model, the company creates a bunch of different agent identities that are each tied to a set of scopes and a whitelist of users. Each agent effectively acts as a unique employee; often, these are owned by individual teams.

So the customer success team may have an agent that only sales and cx folks can talk to, that can access Salesforce but can't access git. And the engineering teams may have an agent that only IT and tech can talk to, that can access git but can't access Salesforce.

src · Agent identity is a really hard problem



This works reasonably well — this is the default at Nori — but it’s a compromise. Large enterprises end up having a bunch of different agent identities, each with slightly different scopes. Also, “ownership” still ends up belonging to the agent, instead of any group of accountable people.

src · Agent identity is a really hard problem



The actual identity layer should live at the conversation level.

The existing auth paradigms are all kinda bad, because they all assume the “agent” is a coherent authentication identity. But it isn’t — every time you call an agent, you may want to do something completely different with it.

Each conversation you have with Claude or Codex should determine what scopes the agent needs for that one conversation. And ideally, there is some avenue for the agent to ask for more scopes as the conversation goes on.

src · Agent identity is a really hard problem



We’re going to be pulling on this thread a lot more.

A lot of this is top of mind because we’re trying to build a more seamless agent identity experience for our Nori cloud infra.

If you have thoughts on how it should work, drop me a line at amol at noriagentic dot com.

src · Agent identity is a really hard problem

