

Agent identity for India: a position paper

Version 0.1 · Published · 12 October 2026

ABSTRACT

AI agents have started to send requests to banks, hospitals and tax offices, and the office that receives one usually has no way to check who sent it or what it may do. This paper describes that problem with invented Indian examples, explains why identity systems that stop at a company's walls leave the gap open, and sets out the principles, roles and headline lifecycle we'd like a shared agent identity for India to follow. It covers how that fits beside W3C credentials, MCP and A2A, what we understand about the Digital Personal Data Protection Act, how we plan to govern the work, and the questions we haven't answered. It is a position, written to be argued with.

READ THIS FIRST

This is a position, not a specification, and not legal advice. The hospital, the GST return and the loan below are cases we made up. Every claim about another organisation or a law comes from a page we read on 12 October 2026, linked where we use it.

Saakshya is a non-profit we're setting up. It isn't incorporated yet, it has no members, and no board. We're being plain about that at the top because a paper about trust ought to start by saying what's missing.

Here is what we'd like the paper to do. We want to say what problem we think India has, in words a person who runs a hospital desk or a loan queue would recognise. We want to say what we'd build on, so you can tell us where we're wrong. And we want to list, near the end, the things we haven't worked out. If you only read one part, read that one.

The paper runs in this order: the problem, why company directories don't solve it, the principles, the roles, a headline-level lifecycle, where this sits next to existing standards, data protection, how we plan to govern the work, and then the open questions. Nothing here describes how the reference implementation works inside. That belongs in an open specification, and we haven't published one.

The problem, through three invented cases

Start with a Monday morning at a hospital.

A request reaches the booking desk at ten past nine. It asks to move a cataract operation from Wednesday to the Tuesday after. The sender says it is an AI agent acting for the patient's daughter, who lives in Dubai and has been trying to arrange the family's calendar around a visit. The patient number is right. The wording is polite. The scheduling system has a rule that says requests from known apps are fine, and this one comes from a known app.

The desk has maybe a minute. Is this really the daughter's agent? Did she ask for exactly this, or did it tidy up a calendar clash on its own? Can it move an operation, or only read the diary? If the patient turns up on Wednesday for a slot that no longer exists, who does the hospital ring? The request can't answer any of that. It only says what it says about itself.

The open question: who should the hospital have been able to ask?

Now a GST return.

A small trader in Surat gives an agent access to the accounting software so it can prepare the monthly return. The trader remembers saying "prepare". The software's permission screen, which the trader clicked through in a hurry, said "file". The agent files. One turnover figure is wrong, and a notice turns up six weeks later. The trader says the agent was only meant to draft. The software company says it did what it was allowed to do. The tax office has a return filed through a portal, with an acknowledgment, and no idea an agent was involved.

All three could be telling the truth. That's what makes it hard.

The open question: where is the record of what the agent was allowed to do, and can anyone outside the trader's own software read it?

And last, a loan.

A comparison agent works for a salaried customer who wants the best rate. In one afternoon it sends applications to three banks. Each bank may pull a credit report, and each sees one application. None of them can see the other two, or tell that all three came from the same agent for the same person. The customer thought they'd agreed to a comparison. Each bank thinks it received a serious application.

The open question: how would a bank know that three requests came from the same agent, for the same person, and what that person agreed to?

These cases look different. The thing they share is plain. An office received a request from software, and it had nothing to look at except the software's own description of itself.

Why a company directory stops at the company wall

It would be unfair to say nobody is working on this. Several large identity vendors now let a company register its own AI agents, and what they've published is serious work.

Microsoft's documentation says Entra Agent ID is generally available (<https://learn.microsoft.com/en-us/entra/agent-id/whats-new-agent-id>). The page carries a 1 May 2026 date and describes owners and sponsors for agent identities, so an organisation can govern its agents the way it governs staff. Okta's post of 29 April 2026 (<https://www.okta.com/blog/ai/okta-for-ai-agents-general-availability/>) says Okta for AI Agents is generally available, and that its Universal Directory is "one place to register AI agents". Ping Identity announced its Identity for AI (<https://press.pingidentity.com/2025-11-06-Ping-Identity-Launches-Identity-for-AI-Solution-to-Power-Innovation-and-Trust-in-the-Agent-Economy>) in November 2025, and a March 2026 release (<https://press.pingidentity.com/2026-03-24-Ping-Identity-Defines-the-Runtime-Identity-Standard-for-Autonomous-AI>) says it is now generally available. We read all four pages on 12 October 2026.

We think that work is good. A company that runs agents should absolutely have it.

But a directory answers questions for the company that runs it. When a bank lists its agent in its own directory, the bank knows who sponsors that agent. The hospital in our first case isn't inside that directory. It's the one acting on the request, and it has no seat at the table. When we read those pages we looked for a documented way for an outside office to check an agent that sits in another company's directory, and we didn't find one. That is a statement about what we read, not about what the products can do. We may have missed something, and we'd like to be told.




Figure 1. Today each company checks only its own agents. A shared register would let them check each other's.

► Read this diagram as text

1. Company A has a wall around its own agents, and trusts only the agents inside it.
2. Company B does the same on the other side.
3. Between the two walls sits a shared register, run by its members and not by any one company.
4. Both companies can look an agent up in the register, so each can check the other's agents.

The picture is the argument. Each company trusts the agents inside its own walls, and nobody can see across. Putting a better directory inside each wall makes the walls tidier. It doesn't put a door in them. We've written the longer version of this in [Why an agent registry has to live outside any one company](https://saakshya.org/articles/an-agent-registry-outside-any-one-company) (saakshya.org/articles/an-agent-registry-outside-any-one-company).

There's a second way to put the same point. Imagine a hospital in Nagpur that gets requests from agents run by a hundred different banks, insurers, apps and individuals. It can't join a hundred directories. It can, in principle, ask one shared place the same few questions about all of them. India has done something like that before. Payments could have stayed a set of bank-owned networks that didn't talk to each other, and a shared rail meant a bank, a wallet and a small shop could check the same thing. We use that only as a comparison. Saakshya has nothing to do with UPI or the bodies that run it.

Platform makers have their own controls too, and they can be careful ones. Our piece on [Meta's Muse agent](https://saakshya.org/articles/muse-can-act-for-you) (saakshya.org/articles/muse-can-act-for-you) goes through why even good controls leave the receiving shop with nothing to check, because they sit inside one company's boundary. That isn't a criticism of the controls. It's a point about who can see them.

What we think should be true

These are our principles. They are the part of this paper we most want argued with, so we've given a reason after each, and said where we're unsure.

The relying party comes first

The relying party is the bank, hospital or office that has to decide whether to act on an agent's request. We start from its desk, because that is where a mistake becomes a real thing: a surgery moved, a return filed, a loan applied for. Many identity schemes are designed around the people who issue credentials, and the person at the front door comes last. We'd like to turn that round.

In practice it means that if a design makes the issuer's life easy and the desk's life hard, we'd pick the desk. It also means the checking has to be quick enough for a person with a queue. We aren't sure yet what "quick enough" is for a small clinic with one receptionist, and we say so again in the open questions.

The person or organisation behind an agent is registered, and their documents are checked, before any identity is issued

An identity is only worth as much as the checking that came before it. If anyone can obtain one by filling in a form, the identity tells an office nothing. So we'd like the order to be: first the individual, company or institution that owns the agent is registered; their documents are checked; the agent's details are checked; and only then is the agent issued an identity.

This is the opposite of how most online accounts work, where you get the account first and the checking, if any, comes later. It's slower at the start. We think it's the right trade, because the whole point is that an office can lean on the result.

What we don't know is how much checking is proportionate. Checking a company's documents is one job. Checking an individual who runs a weekend project is another. We haven't decided where the line sits, and we don't think it should be decided by us alone.

Agents are checked offline, through an SDK

The office that receives a request should be able to check the agent's identity on its own side, using a software kit, without having to depend on a call to someone else at the moment of the decision. A hospital desk on a slow connection, or a branch office in a district with patchy coverage, shouldn't have to go without the check.

That is as far as this paper goes on the subject. How the check works is for the specification, and we won't describe it here.

The specification is open

Anyone should be able to read the rules, implement them, and tell us what's wrong with them. We intend the specification to be royalty-free. A licence saying so is drafted, but not signed, and we won't call it done until it is. A reference implementation exists. The specification itself isn't published yet, and we'll say on this site the day it is.

We care about this for a boring reason. If the rules are closed, an office is trusting a company, and the company is trusting itself. That is the position we're trying to get out of.

India first, and checkable abroad

We're starting in India because the desks we keep picturing, the hospital's, the loan officer's and the tax clerk's, are Indian, and because Indian law is where we'd have to be careful first. But an agent acting for a family in Dubai will send requests to a hospital in Pune, and an agent registered in India will, sooner or later, send one to an office in another country. We'd like an identity issued here to be something an office abroad can check, and we'd like to be able to check agents from elsewhere.

We haven't talked to any foreign body about this. We don't know which countries would want it or what they'd need. It's a direction, not a plan.

Share only what is needed

An office checking an agent should learn what it needs to decide, and no more. A tax office doesn't need to know a patient's hospital. A hospital doesn't need a trader's turnover. Every extra fact that travels with an agent is another fact that can leak, and another thing a data

protection regulator may ask about.

We hold this as a principle and not a feature, because it's easy to agree with and hard to keep when a customer asks for "just one more field". We expect to be tested on it.

Who does what

There are three roles, plus the agent itself. We've kept them few on purpose, because every extra role is another place for blame to slip between.

The owner. This is the individual, company or institution that owns an agent and answers for it. Not "the AI", and not "the platform" unless the platform is the owner. If a pharmacy chain sends an agent to a hospital, the chain is the owner. If an individual sets one up to manage their own paperwork, the individual is. The owner is the person you could write to, and who agreed beforehand to answer for what the agent does.

The registry. This is the role Saakshya is being set up to fill. The registry registers the owner, checks their documents and the agent's details, and issues the agent an identity. It doesn't run the agent, and it doesn't decide what an office should do with a request. It vouches for who stands behind the agent, and for nothing else.

The relying party. The office that receives the request and decides what to do with it. The decision is the relying party's. The registry can tell the desk who is behind an agent. It can't tell the desk to say yes.

The agent is software. It presents its identity along with its requests. As we read Indian law today, an agent isn't a person. The Digital Personal Data Protection Act, 2023 (<https://egazette.gov.in/WriteReadData/2023/248045.pdf>) defines "person" in section 2(s) with a list that includes individuals, companies and firms, and nothing in it names software. We think that means a human or an organisation has to answer, which is the whole reason for the owner role.

There's a question hiding in the word "owner" that we haven't settled. Many agents are built by one company and run by another for a third party's benefit. A bank might take an agent from a vendor and set it loose on its customers' behalf. Who is the owner then, the bank, the vendor or the customer? We lean towards "whoever sent it out and would be answerable to the office", which here is the bank. We expect real cases to be messier than that, and we'd like to hear them.

The life of an agent's identity, at headline level

We'll keep this deliberately short. It has five steps, and the diagram below shows the three questions an office is trying to answer at the end of them.

First, the owner is registered. They tell the registry who they are, and what kind of owner they are: an individual, a company, or an institution such as a hospital or a school.

Second, documents and details are checked. The registry looks at the owner's documents, and at the details of the agent: what it's for, and who runs it day to day. If something doesn't add up, there's no identity.

Third, the identity is issued. The agent now has something an office can check, and the owner has put their name against it.

Fourth, the agent presents it. When the agent sends a request to a bank, a hospital or a tax office, the identity goes with the request.

Fifth, the office checks it. The relying party checks the agent's identity, offline, through an SDK, and decides what to do. We've written a separate paper on what a good front desk should check, which is [a relying party's checklist for AI agents](https://saakshya.org/papers/relying-party-checklist) (saakshya.org/papers/relying-party-checklist).




Figure 2. What an office wants to know before it lets an agent in.

► Read this diagram as text

1. An agent, acting for a person or a company, sends a request to a bank, a hospital or a government office.
2. Before letting the request in, the office looks the agent up in the register.
3. First question: is the agent real? The register shows that it is a registered agent.
4. Second question: who sent it? A named person or company stands behind the agent, and the register says who.
5. Third question: what did it do? There is a record, so anyone can look at it later if something goes wrong.

The diagram is a picture of the three questions, not of how the check is made. Is it real? Who sent it? And is there a record to look at later? The first two are answered by the identity. The third depends on what the relying party writes down for itself, and we think that is the right place for it. The office's record is the office's own.

What this lifecycle leaves out is deliberate. Anything about how the identity is built, how big it is or how it's protected would belong in the specification.

Where this sits next to existing standards

Agent identity isn't a standalone idea. Three sets of standards are close by, and it's worth being clear about what each one does and doesn't cover.




Figure 3. Where Saakshya sits among the standards it is meant to work with.

► Read this diagram as text

1. At the bottom are W3C's open standards for credentials and identity, known as VC and DID.
2. Saakshya sits in the middle, and we are designing it to work with those standards.
3. On one side is MCP, the protocol agents use to reach tools and data.
4. On the other side is A2A, the protocol agents use to talk to each other.
5. Saakshya sits alongside both of them and doesn't compete with either.

W3C verifiable credentials. The W3C published the Verifiable Credentials Data Model v2.0 (<https://www.w3.org/TR/vc-data-model-2.0/>) as a Recommendation on 15 May 2025 (we read it on 12 October 2026). It describes a way to express claims made by an issuer, and it names three parties: an issuer, a holder, and a verifier. Roughly, the registry plays the issuer's part, the agent that presents its identity plays the holder's, and the relying party plays the verifier's. We want to stay compatible with that work, because a shared vocabulary makes it easier for banks and hospitals that already deal with W3C-style credentials. We haven't published how, and this paper doesn't say. We also haven't settled every detail, and we'd rather say that than imply a finished match.

MCP. The Model Context Protocol describes itself as an open-source standard for connecting AI applications to external systems (<https://modelcontextprotocol.io/docs/getting-started/intro>), meaning tools, data sources and workflows. It answers how an agent reaches a tool. It doesn't try to answer who the agent is, in a way an outside office can check. Those are different questions, and the two fit together rather than compete.

A2A. The Agent2Agent protocol describes itself as an open protocol enabling communication and interoperability between agentic applications (<https://a2a-protocol.org/latest/>). The Agentic AI Foundation, hosted by the Linux Foundation, announced on 17 August 2026 (<https://aaif.io/blog/a2a-joins-aaif>) that A2A had joined it as a hosted project, next to MCP. Again, this covers how agents talk. Whether the agent on the other end is who it says it is remains a separate matter. We've looked at this split in Interoperability is standardising. Retention is not. (saakshya.org/articles/interoperability-is-standardising-retention-is-not), and we go through the three standards one by one in Where Saakshya sits next to W3C credentials, MCP and A2A (saakshya.org/articles/where-saakshya-sits-with-w3c-mcp-and-a2a)

So the picture is a layering. Standards for credentials sit underneath. Protocols for tools and for agent-to-agent talk sit beside. An agent identity, as we mean it, tells the office at the end who is behind the agent, in a way it can check for itself. It's meant to work with all three and to replace none of them.

We haven't tested any of this against a real MCP or A2A deployment. The fit, as we describe it, is a design intention and not a result.

We'd also like to be clear about what we've left out. There are other efforts on agent identity, including proposals from researchers that we haven't read closely enough to describe fairly. We'd rather say nothing than describe them from a summary.

Data protection, at a high level

India's Digital Personal Data Protection Act, 2023 is the law we have most in mind. This section is a basis for discussion with counsel. Our separate paper on delegation and accountability under the Act (saakshya.org/papers/delegation-and-accountability-under-the-dpdp-act) goes into it in more detail.

We read the Act and the Rules from the official Gazette copies, and we link to them where we use them. Here are the points that matter most for agent identity.

The Act defines a data fiduciary as "any person who alone or in conjunction with other persons determines the purpose and means of processing of personal data". A hospital that receives a patient's request is one. Section 8(1) says a data fiduciary is responsible for complying with the Act "irrespective of any agreement to the contrary", for processing done by it or on its behalf by a data processor. We read that to mean an office can't hand its duties to an agent, or to anyone, by contract. The Act is in the Gazette of 11 August 2023 (<https://egazette.gov.in/WriteReadData/2023/248045.pdf>), which we read on 12 October 2026.

Second, timing. A commencement notification dated 13 November 2025 (<https://egazette.gov.in/WriteReadData/2025/267647.pdf>) brings most of the Act into force eighteen months after publication, and the provision on registering consent managers a year after it. On the Gazette's dateline that makes 13 May 2027 for most of it. The government's [press release](https://www.pib.gov.in/PressReleasePage.aspx?PRID=2190014) (<https://www.pib.gov.in/PressReleasePage.aspx?PRID=2190014>) says the Rules were notified on 14 November, so the dates may fall a day later. So the main duties aren't running yet, but they are scheduled, and an office adopting an agent check now is building toward them.

Third, what the Act doesn't say. As far as we can tell, neither the Act nor the Rules uses the word "agent" in the sense this paper does. We don't think that is an oversight. The law was written for people and organisations, and it describes duties by role. How those roles apply when software acts for a person is exactly the question we'd like to work on with lawyers and regulators.

We've already written about four ways this shows up. A record that leaves a phone should carry a reference to who authorised it, as we argue in [Local processing is not local authority](https://saakshya.org/articles/local-processing-is-not-local-authority) (saakshya.org/articles/local-processing-is-not-local-authority). A person's right to a summary of who their data was shared with is hard to honour if you can't list the agents that touched it, which [Read-only is a permission](https://saakshya.org/articles/read-only-is-a-permission) (saakshya.org/articles/read-only-is-a-permission) walks through. Where a hosted runtime stores conversation data is a contract detail that no agent protocol settles, as [the retention piece](https://saakshya.org/articles/interoperability-is-standardising-retention-is-not) (saakshya.org/articles/interoperability-is-standardising-retention-is-not) explains. And people who never signed up for an agent can still be affected by other people's agents, which is the worry in our [Muse piece](https://saakshya.org/articles/muse-can-act-for-you) (saakshya.org/articles/muse-can-act-for-you).

None of these is legal advice, and we'd welcome correction.

How we plan to run it

We'll keep this short, because the [governance page](https://saakshya.org/governance) (saakshya.org/governance) says it in full and everything on it is a plan.

The foundation is to be a Section 8 company under India's Companies Act, 2013, called the Saakshya Bharat Foundation. It's in formation. A Section 8 company exists for a public or charitable aim and can't pay dividends to its members. We've chosen that form because the register we're describing is more useful if no single company owns it, and because the rules for it need a home that its users can help steer.

Four kinds of board seat are planned: government, industry, academia, and members elected by the membership. We haven't named anyone, because there is nobody to name. No seat is confirmed. How many seats each kind gets isn't decided either.




Figure 4. How we plan to run it. None of this exists yet.

► Read this diagram as text

1. At the top is the Saakshya Bharat Foundation, a Section 8 non-profit company. It is planned and not yet incorporated.
2. Beneath it are three planned parts: members, working groups and board seats.
3. Members are the organisations and individuals who join. This is planned.
4. Working groups are where members write the standard together. This is planned.
5. Board seats are planned for government, industry and academic representatives.

Four working groups are planned, and none exists yet. One would look at the format of an agent's identity. One would write guidance for relying parties. One would look at how all of it lines up with the data protection law, and would want lawyers in it. One would find words for these ideas in Hindi and other Indian languages, words that people actually use.

At first the foundation runs on grants and donations. We intend the specification to be royalty-free, and a licence saying so is drafted but not signed.

Until a board exists, decisions get made in the open and the reasons get written down. That's a weak arrangement and we know it. It's meant to be temporary.

Saakshya is independent of NPCI, UIDAI, MeitY, RBI, W3C and every company named in this paper. None of them has endorsed it, and we haven't discussed this paper with any of them.

Questions we have not answered

We'd rather list these than hide them. Some are small and some could change the whole plan.

Open questions about the design

What counts as enough checking. We've said an owner's documents are checked before an identity is issued. We haven't decided how much checking suits an individual, a small firm and a large institution, and we suspect the answer differs by sector. A hospital may reasonably want more than a food delivery app.

Who is the owner when several parties are involved. An agent built by one company, run by another and used by a third raises a question we've only answered in the simple case. We'd like examples from people who deal with this every day.

How a small office does any of this. A clinic with one receptionist and no IT department needs a check that takes her seconds. If the check only suits a bank with a security team, the benefit will go to large offices and the rest will carry on trusting names in emails. We haven't solved that.

How granular "what may it do" should be. Permission to book hospital slots is too coarse. Permission to move one particular slot is too fine to manage. Something in between would work, and we haven't found it. A reasonable sector-by-sector answer may be the best we can do.

What happens at the edges of India. We've said "checkable abroad", and haven't spoken to anyone abroad. We don't know how other countries would see an identity issued here, or what we'd need to accept theirs.

What we don't know about people, law and other products

Whether a member-run non-profit stays neutral. Once real regulators and large organisations are involved, the pressure on a small foundation will be considerable. We have a plan for how to try, and we expect it to change when people who've run such bodies read it.

Whether offices will check at all. A desk with a queue can be tempted to accept the request and hope. Whether the checking becomes routine depends on how easy it is and on what regulators expect. Neither is up to us.

How the data protection law will treat agents. We've described our reading, and flagged that the Act doesn't mention agents. Regulators and the Data Protection Board may see it differently. We don't know.

Whether we've got other organisations' products right. We described a few from their own pages, as they stood on 12 October 2026. Those products move quickly. We'd correct any mistake and say that we had.

How to respond

Write to contact@saakshya.org. A reply within a week is what we aim for, and we can't promise more than that yet.

We'd especially like to hear from three kinds of reader. If you run a front door at a bank, hospital or government office, tell us which of our cases is closest to something you've seen, and which one we've got wrong. If you work on identity or credential standards, tell us where we've put the line in the wrong place. If you're a lawyer or work for a regulator, tell us which of our readings of the Act you'd change.

This is version 0.1, dated 12 October 2026. When we change our mind, the old version stays up with a note saying what changed and why. Citing the paper by number and version is the best way to be sure which one you mean.

How to cite this paper

PLAIN TEXT

Saakshya (2026). Agent identity for India: a position paper (Version 0.1). Saakshya Position Paper 1. <https://saakshya.org/papers/agent-identity-for-india>

BIBTEX

```
@techreport{saakshya_agent_identity_for_india_v0_1,  
  title      = {Agent identity for India: a position paper},  
  author     = {Saakshya},  
  institution = {Saakshya},  
  year       = {2026},  
  number     = {Position Paper 1},  
  type       = {Position paper},  
  version    = {0.1},  
  url        = {https://saakshya.org/papers/agent-identity-for-india},  
}
```

Version history

Version 0.1 · 12 October 2026

First version.
