# Who’s the Agent? Building the Identit Layer Healthcare AI Actually Needs

FEB 19, 2026 • PAID



## Table of Contents



## Abstract

This essay argues that agentic AI systems in healthcare require a purpose-built identity and access management layer that doesn’t exist yet, and that building it represents a generational infrastructure opportunity.

### Key points:

- Most agentic AI today piggybacks on user-level credentials, which worked fine agents were glorified macros but breaks badly as they become autonomous actor across multi-system healthcare environments

- Healthcare's regulatory surface area (HIPAA, 42 CFR Part 2, state-level privacy payer contract terms) makes generic enterprise identity solutions like Okta exte non-starters without significant vertical customization

- An agent identity platform for healthcare needs to solve for: fine-grained scopi PHI access, audit trails that satisfy both HIPAA and clinical workflow requirem delegation hierarchies across human and agent principals, and the ability to revo sandbox without breaking workflows

- Market entry probably runs through EHR vendors, health system IT departme healthcare AI middleware companies, not direct to enterprises

- Rough TAM math: ~6,200 US hospitals, 200K+ physician group practices, doze health plans, hundreds of payer-adjacent tech vendors all needing this layer. Conservative ARPU of $50-200K/yr puts addressable revenue in the multi-billio range before you get to international or adjacent verticals

- This is a platform play disguised as a developer tool, and the founders who win understand both OAuth 2.0 scopes and CMS interoperability rules

# The Setup: Why Agent Identity Is a Different Problem Than User Identity

Aaron Levie's tweet is worth sitting with for a minute. The Box CEO isn't usuall guy writing the spiciest technical takes, but he's onto something real here. The r that's governed software authentication for the past 20 years is fundamentally a human-centric one. A person has credentials. A system verifies those credentials system then grants access to things that person is allowed to see and do. OAuth, SAML, SCIM, all of it basically assumes a human is either directly initiating an a or has directly authorized a machine to act on their behalf in a very narrow, pre-defined way.

That model worked beautifully when "agentic" meant a scheduled job that pulle CSV and emailed a report. It's already starting to crack under the weight of wha

agentic AI actually does now, and it'll be completely inadequate within 18 montl

Here's the structural problem. Modern LLM-based agents aren't executing deterministic scripts. They're making judgment calls about what actions to take on context that can shift mid-task. An agent that starts out looking up a patient' medication history might end up writing a prior auth letter, querying a formular database, drafting a message to the prescribing physician, and logging an encour note, all as part of completing one user intent. If that agent is operating on user- credentials, it has access to everything that user can see. And if the user is a hos administrator or a population health analyst, that's a lot of PHI touching a lot of systems for what should be a narrowly scoped workflow.

The blast radius problem Levie mentions isn't hypothetical. It's already happeni non-healthcare contexts, and healthcare is where the consequences of getting it graduate from "bad press" to "federal investigation, OCR audit, and potential criminal liability." The question isn't whether agent identity becomes a critical infrastructure problem. It's who builds the solution first.

## Healthcare Makes This Harder (Obviously)

Let's be honest that healthcare is structurally adversarial to clean technical solut Any founder who's tried to build in this space has a collection of scars from regu complexity, legacy system integrations, and the unique dysfunction of procurem cycles that involve 14 stakeholders, a committee, and a six-month pilot before an signs anything.

Agent identity in healthcare has at least three layers of complexity that don't exi general enterprise software.

The first is the regulatory surface area. HIPAA's minimum necessary standard is essentially a mandate for fine-grained access control. An agent accessing PHI is supposed to access only what's required for the specific purpose. The problem is purpose changes dynamically as an agent executes a multi-step workflow, and HI

enforcement framework wasn't written with agentic systems in mind. The cover entity is still on the hook. The BAA still needs to exist. And the audit logs that C wants to see in an investigation need to tell a coherent story about who accessed and why. That story gets really complicated when the actor is an AI agent that m judgment calls along the way.

42 CFR Part 2 adds another wrinkle for anything touching substance use disord records. These aren't just more restrictive than standard HIPAA controls, they h specific consent and redisclosure rules that need to be enforced at the data layer just at the application layer. An agent that can access a patient's longitudinal rec needs to know, in real time, which parts of that record are Part 2-protected and handle them accordingly. That's not a policy problem. That's an identity and acc architecture problem.

State-level privacy laws are the third dimension. California's CMIA, Texas Healt Safety Code Chapter 181, New York's SHIELD Act and forthcoming health data provisions, all of these layer on top of HIPAA with varying degrees of additional restriction. A health system operating across multiple states needs their agentic systems to enforce jurisdiction-specific rules dynamically. This isn't theoretical. the world as of 2025 and getting more complex, not less.

Then there's the interoperability layer. FHIR R4 and the CMS interoperability ru that went into effect starting in 2021 created a new world where health data is theoretically accessible via standardized APIs. In practice, EHR vendors have implemented SMART on FHIR with varying degrees of fidelity, and the authoriz model that SMART on FHIR specifies was designed for patient-facing apps and clinician-facing apps, not for autonomous agents. The SMART app launch frame assumes a human is in the loop to authorize the OAuth flow. What happens whe there's no human in the loop? That's a design gap that nobody has cleanly solved

## What an Agentic Identity Platform Actually Looks Like

Let's get specific about what this product needs to do, because "agent identity management" is vague enough to mean almost anything and specific enough to n almost nothing depending on who you ask.

At the core, the platform needs to manage what could be called agent principals. These are identities that belong to agents, not users. An agent principal has a de scope of access that's scoped to a specific task type, a specific data domain, or a specific set of systems. It can be provisioned by a human principal who has appropriate authority. It can be delegated permissions up to but not exceeding t delegating human's permissions, which is the same constraint that governs servi account provisioning in traditional IAM but needs new machinery for dynamic contexts.

The scope management piece is where most of the interesting technical work liv Healthcare data isn't flat. A patient record has sections that carry different regu treatments, different clinical sensitivity levels, and different consent states. An a that's doing prior authorization work needs access to diagnosis codes, medicatio lists, and maybe relevant clinical notes, but it doesn't need access to behavioral h records, genetics data, or the patient's billing history. The platform needs to enf that scoping at query time, not just at provisioning time, because the agent can't trusted to self-limit.

Audit logging needs to be a first-class feature, not an afterthought. Every action agent takes, every data element it accesses, every system call it makes, needs to b logged in a format that supports both real-time monitoring and retrospective investigation. The audit trail needs to answer the questions that an OCR investig or a plaintiff's attorney would ask. What did the agent access? Why? Who autho the agent to have that access? What did the agent do with the data? The logging architecture needs to be tamper-evident and needs to support retention policies meet HIPAA's six-year documentation requirement.

Revocation and sandboxing deserve a section of their own. One of Levie's key po about blast radius containment. In healthcare, the ability to revoke an agent's ac or sandbox it without breaking dependent workflows is a genuine operational

requirement. If an agent starts behaving unexpectedly, whether that's a model degradation issue, a prompt injection attack, or just a workflow edge case that w anticipated, a health system needs to be able to pull the agent out of production without taking down the care delivery processes that depend on it. That require identity platform to support graceful degradation, fallback authorization paths, circuit breaker patterns that most enterprise IAM vendors don't offer.

The delegation hierarchy is another piece that's genuinely novel. In traditional I delegation is mostly a solved problem. Users delegate to service accounts, servic accounts have defined scopes, and everything is static. In agentic systems, delega can be dynamic. An orchestrator agent might spin up sub-agents to handle speci tasks and need to delegate a subset of its permissions to those sub-agents. The platform needs to track that delegation chain, enforce that sub-agents can't exce orchestrator's permissions, and maintain the audit trail across the full hierarchy is technically non-trivial and there's no good off-the-shelf solution for it in any vertical, let alone healthcare.

## The Auth Stack: AuthN, AuthZ, and the New Middle Layer

Traditional identity architecture has two main layers. Authentication, which ver who you are, and authorization, which determines what you're allowed to do. Fo human users in enterprise software, this is well solved. SAML handles federated authentication, OAuth 2.0 handles delegated authorization, RBAC or ABAC han permission management, and you're done.

Agentic systems in healthcare need a third layer sitting between AuthN and Aut Call it context-aware access mediation. This is the layer that answers the questic given who this agent is, what it's trying to do, the current state of the workflow i executing, and the regulatory requirements that apply to the specific data it's requesting, what exactly should it be allowed to access right now?

That question is dynamic in a way that traditional AuthZ isn't. RBAC says "this has access to this resource." ABAC says "this attribute combination permits this

action on this resource." Context-aware access mediation says "this agent, execu this workflow step, at this moment in time, acting on behalf of this user, seeking access this specific PHI subset, is or is not permitted to do so, and here's the sco credential that authorizes exactly that action and nothing more." That's a fundamentally different computational problem.

The technical approach that makes the most sense for this is something like a pc engine layered on top of existing standards. Open Policy Agent is the current fav for policy-as-code in cloud-native environments, and it's a reasonable foundatio OPA doesn't natively understand healthcare data models, FHIR resource types, c consent management complexity of multi-state PHI. The value-add is in the healthcare-specific policy library and the FHIR-aware context extraction that fee the policy engine with the right inputs.

SMART on FHIR v2 is worth mentioning here because it does gesture toward so this. The SMART app launch framework updated in 2021 introduced more gran scopes and some support for backend services that don't require interactive use authorization. This is closer to what agentic systems need but still doesn't fully address dynamic scope adjustment mid-workflow or the delegation hierarchy problem. A healthcare agent identity platform that builds on SMART on FHIR v rather than replacing it will have a much easier time with EHR vendor integratic because those vendors have already built SMART support into their API layers.

The authentication piece for agents is actually simpler than the authorization pi but it's not trivial. The current best practice is mutual TLS with short-lived certificates, combined with a machine identity provider that can attest to the age software composition. This is standard in cloud-native infrastructure, less stand healthcare IT. Many health system environments still have legacy systems that dc support mTLS, and getting those systems to participate in a modern agent ident framework requires either API gateway intermediation or protocol translation la A platform that handles this gracefully, that can issue agent credentials that wor across both modern FHIR APIs and legacy HL7 v2 interfaces, has a real advantag

# The Market Opportunity and Go-To-Market Logic

The TAM arithmetic here is interesting. US hospitals alone number around 6,20 Physician group practices at scale number north of 200,000, though consolidatic reducing that number and concentrating buying power in fewer hands. Health p PBMs, and their tech vendors add another layer. If the pricing model is infrastru software, the right analog is something like Okta or CrowdStrike, where you're s a platform that's embedded in critical operations and commands pricing commensurate with the risk it mitigates.

Conservative assumption: a health system with 500+ beds is paying $100-200K p year for an agent identity platform once it has meaningful AI deployment. There roughly 2,000 US hospitals in that size range. That's $200-400M in ARR just fro large hospitals before you touch ambulatory, health plan, or health tech vendor segments. The health tech vendor segment is actually the more interesting near- opportunity because those companies are the ones deploying agentic AI fastest a they need an identity layer they can embed in their products rather than asking t customers to procure separately.

Go-to-market for infrastructure like this almost always runs through a bottlenec security, it runs through CISOs. In identity, it often runs through IT or engineer leadership. In healthcare, there's an additional bottleneck which is whoever own EHR relationship, because every other system integration flows through it. That means partnerships with Epic, Oracle Health, athenahealth, and the other major vendors are strategic, not just tactical. An agent identity platform that's listed in Epic's App Orchard and certified against Epic's FHIR APIs has a fundamentally different sales motion than one that requires custom integration work at every h system.

The venture-backed healthcare AI companies are the fastest path to initial reven Companies building clinical documentation tools, prior auth automation, care g identification, population health management, all of them are deploying agentic workflows and none of them have solved the identity problem cleanly. Selling to

companies as a developer tool, with an SDK and a managed service layer, gets yo distribution through their deployments without having to do direct health syste sales. The tradeoff is margin compression and customer concentration risk, but early-stage company it's the right trade.

## Risks, Moats, and Why This Won't Be Easy

The competitive risks are worth being honest about. The biggest one is that the hyperscalers solve this themselves. Azure Health Data Services already does som what's described here. AWS HealthLake has identity and access controls. Google Cloud Healthcare API has similar capabilities. If Microsoft or Amazon decides t healthcare agent identity is a priority and integrates it deeply into their existing healthcare cloud offerings, a standalone platform is competing against free-as-in bundled.

The counter to that is that healthcare IT is genuinely heterogeneous in a way tha makes single-cloud solutions rare. Most large health systems are multi-cloud or prem plus one cloud, and they're not going to standardize their agent identity infrastructure on a single hyperscaler. The standalone platform that works acros cloud environments and on-premises systems has a structural advantage that the hyperscaler solutions can't easily replicate.

The EHR vendors themselves are another competitive risk. Epic has enough mar share to build something like this and make it the default for all Epic customers have the relationships, the data models, and the existing integration surface area risk here is lower than it might seem because Epic tends to move slowly on new infrastructure categories and because health systems have real appetite for solut that reduce their dependence on Epic's roadmap decisions. But it's a risk worth modeling.

The moat for a standalone platform in this space comes from a few places. Regul expertise compounded into the product is probably the strongest one. If the poli engine has a well-maintained, continuously updated library of healthcare-specifi

authorization rules that reflect current HIPAA guidance, OCR enforcement tren and state-level privacy law changes, that's genuinely hard to replicate quickly. T second moat is network effects in the audit and benchmarking layer. If enough h systems are using the platform, the aggregate audit data becomes a source of ins about normal and anomalous agent behavior that you can feed back into detectic policy recommendations. That's a data flywheel that a new entrant can't access.

Customer switching costs are another structural moat. Identity infrastructure, o embedded, is genuinely painful to change. The health system that builds their ag deployment architecture around this platform's identity primitives is not casuall evaluating alternatives a year later. That's not to say switching is impossible but inertia is real and valuable.

## What Founders Should Build Right Now

The actionable framing for founders thinking about this space is to start with th audit problem because it has the shortest path to revenue and the clearest buyer Health systems already know they need better audit trails for AI systems. OCR h yet issued specific guidance on AI agent access logs but every health system CIS anticipating that it's coming and wants to be ahead of it. A product that solves specifically for AI agent audit logging, that produces HIPAA-compliant audit rec for agent actions in a format that's interpretable by compliance teams and ready OCR review, can be sold today without requiring the full identity platform build

From there, the expansion path is natural. Once you're in the audit layer, you ha visibility into what agents are doing and what access they're using. That visibilit the foundation for scope recommendations, which is the foundation for scope enforcement, which is the full identity platform. It's a land-and-expand motion t doesn't require a health system to bet their entire AI infrastructure on a new ver before they trust you.

On the technical architecture side, the non-negotiable early decisions are aroun policy engine design and the FHIR integration approach. Building on OPA gives strong foundation with active community development and good enterprise tool

Building FHIR integration on SMART on FHIR v2 rather than proprietary APIs you out of vendor lock-in risk and makes EHR certifications tractable. These cho constrain your near-term flexibility but dramatically reduce your integration bu as you scale.

The team composition matters a lot here. This is genuinely a domain where tech depth and regulatory expertise need to coexist in the founding team. A team tha engineers will build something technically elegant that doesn't account for the nuances of HIPAA enforcement patterns or the political dynamics of health syst procurement. A team that's all healthcare domain experts will understand the problem deeply but will struggle to make the right technical architecture decisi The sweet spot is probably a CTO-level person with cloud identity infrastructur experience, a product or GTM lead with health system or health tech company experience, and early access to healthcare regulatory counsel who's embedded er in the product to influence design decisions rather than just reviewing output.

The timing argument for building this now rather than in two years is straightforward. The agentic AI deployments are happening now. Health system running pilots and early production deployments of AI agents without having so the identity problem. They're making do with user-level credentials and manual access reviews and they know it's not sustainable. The regulatory guidance is goi come and when it does, the companies that are already in production with a solu will have a structural advantage that late entrants can't easily overcome. The win for being the default infrastructure layer for healthcare agent identity is open ri now and it probably closes within 24 months as the larger players catch up or as early mover establishes enough market position to make the category effectively closed.

Healthcare AI is real this time, and the infrastructure layer that enables it to be deployed safely and compliantly is one of the most interesting unsexy bets in the market. Agent identity isn't the flashiest pitch but it's the kind of platform that up being worth more than most of the AI applications it enables.

2 Likes • 2 Restacks

← Previous

Nex

## Discussion about this post

Comments Restacks

Write a comment...

Jim StClair Feb 19

Holey Moley, you crawled inside my head and pulled out what I've been working c
a significant team of volunteers across the spectrum). Well done!

LIKE REPLY