

The HIMSS Conference Nobody Actually Attended

MAR 11, 2026 • PAID



Share

Abstract

The conference floor at HIMSS is already kind of a simulation. Anyone who has walked the Las Vegas Convention Center in March, badge scanning through vendor booths while someone in a logo polo tries to schedule a “brief 30-minute follow-up” knows the theater of it. The whole thing is performative sales infrastructure dressed up as knowledge transfer. Which makes the following thought experiment less cynical than it sounds: what if the humans just stopped going, and sent their AI agents instead?

This is not a piece about the metaverse. That hype cycle is mostly dead and nobody is pouring another drink for it. This is about something that is actually being built now, across a dozen enterprise software stacks simultaneously, and the question is whether the health tech industry is going to be intentional about it or stumble into it backwards.

Key Arguments:

- The primary value of conferences like HIMSS is top-of-funnel discovery and partnership exploration, not the content sessions
- AI agents are already capable of executing the semantic matching, negotiation scaffolding, and follow-up workflows that conferences are supposed to accelerate
- A new business model exists for conference operators who want to survive the transition rather than be bypassed by it

- The technical architecture is largely available today using existing LLM infrastructure, RAG systems, and multi-agent orchestration frameworks
- The human role shifts from “attendee” to “deal closer” with significant product implications for enterprise health tech sales teams

Data Points Referenced: HIMSS 2024 attendance (~28,000 registered), average enterprise sales cycle in health IT (12-18 months), cost per attended lead at major health tech conferences (\$800-\$2,000+ all-in per meeting), LLM context window capabilities relevant to negotiation tasks.

Table of Contents

The Conference as a Broken Sales Primitive

What AI Agents Can Actually Do Today

The Architecture of an Agent-to-Agent Conference

Business Models for the Operator and the Attendee

The Human Closes the Deal

Why HIMSS Should Build This Before Someone Else Does

The Conference as a Broken Sales Primitive

Here is what HIMSS actually is if you strip away the educational programming, the keynotes, the Harold Klieman Foundation awards, and the slightly depressing business situation: it is a very expensive, very slow, very geographically concentrated CRM event. Companies spend north of \$100,000 just to have a booth presence when you factor in the exhibit space, shipping, booth build, hotel rooms, flights, meals, and the opportunity cost of pulling your best people off their actual jobs for a week. In reality, they get badge scans, some conversations that were already scheduled two months

advance anyway, and a pile of business cards that someone will fail to enter into Salesforce before February is over.

The content sessions are largely a sideshow. The real work of HIMSS happens in Starbucks line, the suite meetings at the Wynn, and the private dinners that the companies with real budgets host on Tuesday night. The conference is functioning as a temporal and geographic forcing function. It gets the right people in the same place at the same time, which is genuinely valuable, but the mechanism is preposterously inefficient relative to the outcome.

To be specific about the numbers: HIMSS 2024 drew approximately 28,000 registered attendees. Of those, a meaningful fraction are there in explicitly commercial roles, meaning they are either selling, buying, or evaluating partnerships. If you assume conservatively that the average enterprise health tech company spends \$150,000 to attend, and that they walk away with 40 qualified conversations, the cost per meaningful interaction runs somewhere between \$800 and \$2,000 when you account for the ones that go nowhere. That is an extraordinarily high price for what is essentially a discovery mechanism, a way to find out that a company you had never heard of has a data asset or a workflow capability that maps onto a problem you are actively trying to solve.

The tragedy is not that conferences exist. The tragedy is that the discovery layer, the part where you figure out who you should be talking to and why, is done almost entirely by humans wandering around a floor plan with a printed map, hoping to bump into the right booth. In what other domain would you accept that level of signal-to-noise for that price?

What AI Agents Can Actually Do Today

The relevant capabilities here are not theoretical. They are deployed at scale in enterprise software contexts right now, just not explicitly in conference simulations.

Modern LLM-based agents can read, parse, and reason over complex organizational documents, product descriptions, capability matrices, and contract structures. V

retrieval-augmented generation, an agent can ingest a company's website, press releases, SEC filings if public, partnership announcements, API documentation, pricing pages, and construct a reasonably accurate model of what that company who it sells to, what problems it solves, and what a partnership with them might like. This is not magic. This is just document QA at scale with some reasoning layer on top.

What is newer, and more interesting for this use case, is the multi-agent coordination layer. Frameworks like LangGraph, AutoGen from Microsoft, and CrewAI allow developers to spin up multiple agents with distinct roles, personas, and objectives and have them interact with each other through structured conversation. An agent representing a revenue cycle management company can be given a system prompt that includes that company's partnership criteria, ideal customer profile, existing integration ecosystem, and business development priorities. Another agent representing an EHR vendor can be given its own equivalent context. The two agents can then have a conversation that explores potential integration points, revenue structures, data licensing opportunities, and co-marketing angles, and produce a structured output summarizing where alignment exists and where it does not.

This is not the same as two chatbots talking to each other about nothing. The quality of the output is a direct function of the quality of the context each agent has been given. If you feed an agent a detailed company brief that includes things like "we have a gross margin profile that makes revenue share models attractive above 15% and unattractive below 12%, and we are actively seeking partnerships with companies that have EHR integrations in the Epic App Orchard," then its outputs will reflect those constraints in ways that are actually useful to a human business development professional reviewing the transcript later.

The agents can also run these conversations in parallel, across dozens or hundreds of potential partners simultaneously, in the time it would take a human BD rep to have one 20-minute conversation at a conference booth. The productivity math here is subtle.

The Architecture of an Agent-to-Agent Conference

Building this as a product requires thinking about several distinct infrastructure layers, and it is worth being specific about what each one does.

The first layer is the company knowledge graph. Every participating company needs to have its agent loaded with accurate, current, structured information about what the company does, what it is looking for, and what constraints exist on the kinds of deals it can pursue. This is harder than it sounds. Most companies do not have this information in a form that is easily digestible by an LLM. It lives in pitch decks, Confluence pages, in the heads of the BD team, in informal tribal knowledge about why the last three partnership conversations fell apart. Building the tooling to extract and structure this information is itself a product problem, and probably the hardest one in this stack. The knowledge graph layer would need to include product capabilities, integration technical requirements, commercial model preferences, compliance constraints like HIPAA BAA requirements or data governance policies, existing partner relationships, and what a good versus bad fit looks like. The companies that have invested in product-led growth tooling, where they have had to articulate their ICP and integration requirements in formal terms, will have a much easier time populating this than companies where partnership strategy lives in a person's head.

The second layer is the agent runtime and persona management. Each company's agent needs to have a stable identity, a set of objectives, and a negotiation posture that reflects the company's actual priorities. This is closer to prompt engineering than software engineering, but at scale it requires systematic tooling. You need to be able to version the agent's persona, update it when the company's priorities change, and audit its behavior to make sure it is not accidentally committing to things the company would not actually agree to. Guardrails matter here. An agent that is too aggressive might generate partnership proposals that the human reviewers immediately discard as unrealistic. An agent that is too conservative produces opportunities that are too vague to be actionable. Tuning this is an empirical process, and the

framework probably involves red-teaming the agents against each other before deploying them in a live environment.

The third layer is the matchmaking and scheduling orchestrator. This is the piece that determines which agents talk to which other agents, and in what order. A naive implementation just runs every agent against every other agent, which is computationally expensive and produces a lot of noise. A smarter implementation adds a semantic similarity layer, probably embedding-based, to pre-screen pairs of companies for potential fit before spinning up a full agent conversation. Think of it as the conference floor layout problem. At a real conference, the exhibitor floor is organized so that buyers walk past relevant sellers. The orchestrator is doing the work algorithmically. It might use a combination of stated partnership preferences, industry taxonomy tags, revenue stage filters, and integration ecosystem membership to generate a ranked list of high-probability pairs. Only the top-ranked pairs get full agent conversation treatment.

The fourth layer is the conversation framework itself. Agent-to-agent conversations need structure to be productive. Left to their own devices, LLMs in a multi-agent setup can produce conversations that are fluent but content-free, full of polite acknowledgment and vague alignment language that sounds like a conference call that nobody wanted to be on. The conversation framework needs to include explicit milestones: an opening exchange where each agent summarizes its company's capabilities and priorities, a structured exploration phase where the agents probe specific fit dimensions, a conflict identification phase where genuine misalignment is surfaced rather than papered over, and a closing synthesis that produces a structured summary of what was found.

The fifth layer is the human review and escalation interface. The agents are not making decisions. They are generating discovery artifacts that humans use to make decisions. The output of each agent-to-agent conversation needs to be presented to the relevant humans in a format that allows rapid triage. Something like: "You talked to 47 companies over the past 72 hours. Here are the 8 conversations where strong potential fit was identified, ranked by estimated deal size and strategic alignment. Here are the 12 where moderate interest exists but one open question

needs a human judgment call. The other 27 are archived.” The human BD team spends their time reviewing and acting on pre-qualified opportunities rather than generating them from scratch.

The sixth layer, and the one that makes this feel like an actual conference rather than just a batch API call, is the asynchronous communication protocol. Real conference conversations are not one-shot exchanges. They are iterative. Someone says something interesting and the other person asks a clarifying question. A detail comes up that changes the frame. The agent conversation framework needs to support multiple turns, probably with some constraint on total length to force prioritization, and needs to handle the situation where one agent surfaces a question it cannot answer from its knowledge graph, which then gets flagged for human input before the conversation continues.

Business Models for the Operator and the Attendee

If you are thinking about building this, there are at least three distinct business models available, and they are not mutually exclusive.

The first is the conference operator model, where an existing player like HIMSS Global itself, or an independent startup, runs the platform as a managed service. Companies pay a registration fee to participate, analogous to exhibit fees today, dramatically lower given the absence of physical infrastructure costs. The platform operator handles agent onboarding, runs the matchmaking orchestrator, manages conversation infrastructure, and delivers outputs. Revenue is a combination of flat registration fees and premium tiers for things like higher conversation volume, priority matching slots, or enhanced analytics on how your agent performed across a session. The economics here are interesting because the marginal cost of adding another participating company is much lower than in a physical conference context, which means the platform can afford to serve smaller companies that would never have the budget for a real HIMSS booth. A startup with a \$2 million seed round cannot afford \$100,000 for HIMSS floor space. They might absolutely pay \$5,000

\$15,000 to have their agent participate in a virtual conference session where they matched against 30 relevant potential partners.

The second model is the enterprise subscription play, where the platform is sold to individual companies as a year-round BD intelligence tool. The conference event is just the highest-density version of something the company is running continuously. Between conference cycles, their agent is passively monitoring the market, scanning new company announcements, flagging potential partners that meet their criteria, and running exploratory conversations with targets that opt into the system. The annual subscription includes participation in quarterly high-density “conference events” and ongoing access to the agent runtime and matchmaking infrastructure. This is a close analog to what intent data vendors like Bombora or G2 sell today, but with a much higher-fidelity output because instead of just telling you that a company is researching your category, it gives you a transcript of what that company’s AI agent said when it talked to yours.

The third model is the deal advisory layer, where humans who are expert in health tech partnerships use the agent conversation outputs as the raw material for a managed service. The platform operator provides both the technology and a team of experienced BD professionals who review the agent outputs, prioritize the opportunities, prepare the humans for follow-up conversations, and track pipeline development over time. This is the model that makes the most sense for earlier-stage companies that have good technology but thin BD infrastructure. You get the discovery efficiency of the agent layer plus the commercial judgment of experienced humans who have actually done health system deals before. The margin structure in this model is more like a consulting or fractional BD firm than a SaaS business, but the unit economics can work if the client base is large enough.

There is also a data business embedded in all three of these models that is worth naming explicitly. The aggregate pattern of what agents say to each other, which partnership proposals get escalated to human review, which deal structures show up repeatedly across conversations, is extraordinarily valuable market intelligence. A platform running thousands of agent conversations per quarter across the health ecosystem would have a real-time picture of where commercial interest is clustered.

which business models are attracting partnership appetite, and which categories companies are struggling to articulate their value proposition clearly enough for agent to generate quality discovery outputs. That intelligence, appropriately anonymized and aggregated, is worth money to investors, analysts, and the companies themselves. The platform operator sitting on top of that data asset needs to handle carefully from a privacy and competitive sensitivity standpoint, but the strategic value is obvious.

The Human Closes the Deal

Nothing in this framework suggests that the humans go away. The humans are the most important part of the system. They are just doing different work.

What the agent-to-agent conference replaces is the expensive, inefficient, attention-intensive work of early-stage discovery. It replaces the wandering-around-the-floor phase, the “let me grab a card and have my team follow up” phase, and a lot of the initial exploratory calls that currently consume enormous amounts of BD bandwidth to produce information that could have been gathered algorithmically. Enterprise health tech sales cycles run 12 to 18 months on average. The discovery phase of the cycle, finding the right counterparty and establishing that there is something worth exploring, can consume two to four months of elapsed time and significant human effort. If an agent can compress that to a week of high-volume parallel conversations, the time saved is commercially significant.

What the agent cannot do is build the trust required to close a complex deal. Health tech partnerships, particularly ones that involve data access, clinical workflow integration, or risk-sharing arrangements, require a level of relationship capital and institutional credibility that cannot be established through agent transcripts. As an example, a CEO at a large health system is not signing a seven-figure data licensing agreement because two AI agents had a productive conversation. What the agent conversation does establish is that there is a credible hypothesis worth testing, that the companies are roughly aligned on value proposition and commercial model, and that there are 1

obvious disqualifying factors. It gets the humans to the table having already done work that usually takes months of cold outreach and preliminary calls.

This changes the job description of a health tech BD professional in a meaningful way. Less time on discovery, more time on qualification and relationship development. Less time on first calls that go nowhere, more time on second and third calls that have real momentum. The BD professionals who thrive in this environment are the ones who are good at the genuinely hard parts of enterprise sales: navigating stakeholder complexity, managing the political dynamics of large health system decisions, reaching the room in ways that no agent currently can, and building the credibility and personal trust that makes a counterpart comfortable committing to a partnership that involves risk.

The productivity implications for enterprise sales teams are real. A BD team of 10 people, currently managing an average of maybe 20 to 30 active pipeline conversations at any one time given the bandwidth constraints of discovery work, could plausibly manage 80 to 100 active conversations if the discovery layer is automated. Not all those conversations become deals, obviously, but the law of large numbers matters in health tech BD the same way it matters everywhere else.

Why HIMSS Should Build This Before Someone Else Does

HIMSS, the organization, has assets that a startup trying to build this from scratch does not have. It has relationships with essentially every major health tech company and health system in the country. It has a trusted brand in a context where trust matters for data sharing. It has a decade of structured data about which companies show up to talk to which other companies, which gives it a massive head start on the matchmaking quality problem. And it has the convening authority that would make it easy to get companies to opt in, because participating in the HIMSS agent network would be a natural extension of an existing relationship rather than a cold sales proposition from an unknown startup.

The business risk HIMSS faces is not that the agent conference replaces the physical conference overnight. It is that a well-funded startup builds the agent conference first, establishes the network effects that come from having the most companies participating, and gradually shifts the center of gravity of health tech commercial activity away from the physical event. By the time HIMSS notices, the startup is *facto* venue for discovery-stage health tech BD, and HIMSS is left running a convention that is increasingly about content delivery rather than commercial activity.

This is the same dynamic that has played out in other conference-dependent industries when digital platforms entered. The physical event often survives, but its role changes. It becomes a relationship-deepening venue for people who already know each other from the digital platform, rather than a discovery venue for people who do not. That is a meaningful degradation in commercial relevance and pricing power for the conference operator.

The obvious strategic play for HIMSS is to build or acquire the agent conference capability and position the physical event as the premium tier of a year-round platform. The agent network runs continuously, the physical event in March is where the relationships established through agent conversations get formalized and deepened. Sponsorship dollars shift toward the platform rather than the floor, the exhibit fee model evolves into something that reflects the year-round value of the commercial network rather than just the week-long event.

None of this requires capabilities that do not exist. The technology is available. The business model math works. The strategic case for the incumbent is clear. What is missing is the organizational will to cannibalize a profitable existing business model before a competitor does it instead. That is, of course, the hardest part of any platform transition in any industry. But health tech has a reasonable track record relative to healthcare delivery itself, of adapting to technology transitions when commercial incentives are sufficiently clear.

The agent conference is coming one way or another. The only real question is who builds it, and whether the companies that have the most to gain from network effects

in health tech commercial activity are the ones running it or the ones paying to ε
it.

[← Previous](#)

[Next →](#)

Discussion about this post

Comments

Restacks



Write a comment...

