Infinitus

To Infinitus & beyond

Building the continuous patient journey

ANKIT JAIN, SEPTEMBER 2026

Table of contents

A thirty-year horizon

People often ask where I hope Infinitus will be five years from now. I've never found that to be a particularly interesting question. Five years is enough time to launch products, enter markets, and build a large company. It isn't enough time to know whether you've built something that truly matters. I think in decades instead. A thirty-year horizon changes how you build because almost everything you can see today will change. The models will be replaced. The interfaces will evolve. Entire technology stacks will disappear. That's never worried me because I've never believed we're building an AI company. I've always believed we're building a healthcare company that happens to use AI. Technology is simply today's best tool for solving a problem that I believe will outlast the tool itself.

The problem we've chosen is one that I've become increasingly convinced sits at the center of modern healthcare. Healthcare runs on conversations, yet those conversations have never had the infrastructure they deserve. As patients move between physicians, payors, specialty pharmacies, manufacturers, and health systems, context disappears, attention becomes fragmented, and waiting becomes an accepted part of care. The people inside those organizations are rarely working against one another. In fact, I've found the opposite. Almost everyone is trying to help patients receive better care. What differs isn't the destination. It's the information each participant has, the constraints they operate under, and the systems they've inherited. Patients experience one continuous journey. The healthcare system rarely does.

This book is my attempt to explain how we've come to see that problem and why we've chosen to spend decades working on it. It begins with the principles that shape how we think, then explores why healthcare feels fragmented despite the good intentions of the people inside it, why this moment in technology is unlike any before it, why we believe communications, not records, will become the next foundational layer of healthcare, and finally the culture and leadership principles we've tried to build while pursuing that mission. None of these chapters are meant to predict exactly what healthcare will look like thirty years from now. They are an attempt to describe the beliefs that I hope will still be true even after today's technology has become tomorrow's history.

If there's one idea I hope readers take away, it's that technology should never become the story. The story is the patient. If we can preserve context across every conversation, expand the amount of attention the healthcare system can provide, and allow clinicians to spend more of their time caring for people instead of coordinating paperwork, then the specific models, architectures, and interfaces we used along the way won't matter very much. They'll simply have been the tools that helped make healthcare feel a little more human.

ACT 0

The Infinitus truths

We believe the biggest problem to solve in healthcare is coordination.

We believe patients shouldn't have to be the glue that holds their healthcare journeys together.

We believe healthcare runs on communication more than it runs on software.

We believe every conversation should make the next conversation better.

We believe the platform should be biased toward the patient, not the automation.

We believe the best agent isn't always AI. The best agent is the one that best serves the patient at that moment.

We believe AI's greatest contribution won't be replacing work, it will be making possible the care that was not delivered because time was scarce.

We believe healthcare is different, and software should respect that difference.

We believe trust is earned in inches.

We believe customers should spend their time building what makes them unique, not reinventing the foundations of healthcare AI.

We believe patients shouldn't adapt to technology, technology should adapt to patients.

We believe technology changes. Patient needs don't.

We believe great companies relentlessly reinvent their products while remaining obsessed with the same customer problem.

ACT I

Why healthcare feels broken

It shouldn't take eight weeks

The neurologist had done the hard part. After weeks of appointments, MRIs, blood tests, and uncertainty, the neurologist finally had an answer. It wasn't the answer anyone hopes for, but it was an answer. The patient had multiple sclerosis. Fortunately, this wasn't 20 years ago. Modern medicine has transformed the outlook for patients with MS. Therapies now exist that can dramatically slow progression, reduce relapses, and preserve quality of life, especially when treatment begins early. The physician explained the options, recommended a therapy, answered every question, and sent the prescription before the patient even left the office. From a medical perspective, everything had gone exactly as it should have.

The next eight weeks had almost nothing to do with medicine.

The prescription arrived at a specialty pharmacy. Benefits were verified. A prior authorization was submitted. Additional clinical documentation was requested from the physician's office. The documentation was sent. The request was reviewed and denied. An appeal was filed. Meanwhile, another problem surfaced. Even if the appeal succeeded, the patient's out-of-pocket responsibility would be thousands of dollars, so applications for copay assistance had to be completed, reviewed, and supplemented with additional information before they could be approved.

None of these steps was unreasonable on its own. Collectively, they became an obstacle course. Every organization involved was doing real work. Every phone call moved one small piece of the process forward. Every approval unlocked the next dependency. Yet from the patient's perspective, nothing seemed to be happening. The only thing the patient could measure was time. Eight weeks after her neurologist decided treatment should begin, the patient finally received her first infusion.

When people hear stories like this, the instinct is to find the villain. Someone must have dropped the ball. Someone must have made the wrong decision. The longer I've spent inside healthcare, the less convincing that explanation has become. The neurologist wanted treatment to start immediately. The specialty pharmacy was doing exactly what it had been asked to do. The manufacturer had invested not only in discovering the therapy but also in building patient support programs to make it affordable. The payor was applying its clinical coverage policies. Every organization believed it was helping the patient, and in many ways every organization was right.

The delay wasn't caused by one bad decision or significant mistake. It emerged from dozens of perfectly reasonable decisions and actions that had to happen in sequence across organizations that were never designed to work as one.

That observation changed the way I think about healthcare.

Healthcare is often described in terms of physicians, hospitals, payors, manufacturers, pharmacies, and countless other participants. Patients don't experience any of those distinctions. They experience one journey. They don't remember which organization requested another form or who processed the authorization. They remember waiting. They remember explaining their situation over and over again. They remember wondering whether anyone actually knew where they were in the process.

I've come to believe that some of healthcare's biggest problems don't exist inside organizations. They exist between them.

Over the last several decades, healthcare has become extraordinarily good at specialization. Physicians diagnose diseases that would have been impossible to recognize a generation ago. Manufacturers continue developing therapies that fundamentally change patients' lives. Specialty pharmacies have built sophisticated support organizations. Payors evaluate millions of clinical decisions every year. Each participant has become remarkably good at its own part of the journey. The journey itself, however, remains fragmented.

One of the reasons is surprisingly simple. Healthcare doesn't move because information exists. It moves because people communicate. A diagnosis is explained through a conversation. A prior authorization is negotiated through a conversation. A care transition depends on a conversation. A specialty medication reaches a patient through dozens of conversations connecting physicians, payors, pharmacies, manufacturers, infusion centers, and patients. Those conversations are not administrative overhead. They are the mechanism through which healthcare operates.

When they become fragmented, care becomes fragmented.

The patient I mentioned above didn't wait because medicine failed. She waited because coordination failed. Those are very different problems, and they lead to very different solutions. If the problem is medicine, we need better drugs. If the problem is coordination, we need a better way for healthcare to communicate.

That's the idea that eventually became Infinitus.

The real bottleneck: coordination and attention

Every generation inherits a different bottleneck. For most of modern medicine, the constraint was knowledge. We didn't understand disease well enough. We couldn't diagnose it early enough. We didn't have effective therapies once we did. The defining breakthroughs of the last century came from expanding what healthcare knew.

Today's bottleneck is different. Medicine already knows an extraordinary amount. We diagnose diseases earlier than ever before. We have therapies that dramatically alter the course of cancer, autoimmune disease, rare disease, and countless chronic conditions.

Yet patients still miss screenings, delay specialist appointments, wait weeks to begin therapy, abandon treatment before it starts, and quietly disappear from care.

The problem increasingly isn't what healthcare knows. It's whether healthcare can consistently apply what it already knows to every patient who needs it.

Across physicians, payors, manufacturers, specialty pharmacies, and health systems, I've found remarkable alignment on the outcome everyone is trying to achieve: the right patient starts the right therapy at the right time with the right support. The problem is that every organization sees only the portion of the journey that passes through its own walls. The patient experiences only the spaces between them.

Healthcare's biggest bottleneck is coordination. Coordination requires attention, and attention has always been scarce. Physicians would spend more time answering questions. Nurses would check on more patients after discharge. Specialty pharmacies would rather call before a refill is missed than after adherence has already become a problem. Care coordinators know exactly which patients would benefit from another conversation. Those conversations rarely fail because they aren't valuable. They fail because there has never been enough capacity to make them possible.

Healthcare has never rationed compassion. It has rationed attention. Scarce attention inevitably becomes waiting. Waiting for a screening. Waiting for a referral. Waiting for blood work. Waiting for prior authorization. Waiting to begin therapy. Waiting for someone to call back. Healthcare measures those delays operationally—days to therapy, turnaround times, service levels. Patients experience them biologically and emotionally. Disease progresses. Anxiety compounds. Questions accumulate.

Questions accumulate.

- *Is this treatment safe?*
- *Can I afford it?*
- *Is anyone actually working on my case?*

The longer those questions go unanswered, the more likely patients are to miss appointments, abandon therapy, or simply disengage from care.

Waiting isn't just an operational metric. It's a clinical outcome. Every unnecessary day between symptom and diagnosis, diagnosis and treatment, treatment and follow-up increases the probability that something goes wrong—not only because disease advances, but because uncertainty changes behavior. Patients don't simply wait. They make decisions while they're waiting.

That's why I think the next generation of innovation in healthcare will look different from the last. The last fifty years were largely about expanding medical knowledge. The next fifty will increasingly be about expanding healthcare's capacity to coordinate, communicate, and deliver that knowledge. Better medicine matters enormously. So does making sure patients actually receive it.

ACT II

Why Infinitus exists

Healthcare runs on conversations

Imagine if every corporate email vanished the moment it was sent, or if every meeting ended with absolute amnesia. No modern business could survive that way, yet that's exactly how healthcare operates. We log the final clinical outcome, but we lose the vital context: the questions asked, the concerns raised, and the promises made along the way. Because conversations disappear the moment they end, every new interaction is forced to start from scratch.

A patient doesn't start treatment because a database updated; they start because a doctor explained a diagnosis, an insurer approved a prior authorization, and a nurse reassured them over the phone. Nearly every critical milestone, including a referral, an appeal, a medication change, and a discharge plan, is fundamentally a conversation.

Healthcare is unique because the core challenge isn't processing transactions; it's coordinating people. Yet, while we've spent billions of dollars digitizing patient records, we've left the communication channels largely untouched. Phone calls still sit on hold, faxes are still flying, and patients are still forced to repeat their life stories to every new organization they encounter. We became world-class at storing information while leaving the communication that moves that information frozen in time.

This is the missing layer. Healthcare doesn't need more software; it needs a way for context to flow seamlessly across the patient journey. Whether an interaction happens over voice, text, email, whether it's with another person or through an AI agent, it should build on the last conversation instead of wiping the slate clean. The future of healthcare won't be defined by replacing human conversations, but by finally making every single one of them count.

The missing layer

Every industry eventually develops foundational infrastructure that quietly makes everything else possible. Banks rely on payment networks. The internet runs on TCP/IP. Enterprises depend on operating systems, databases, and cloud platforms. Healthcare built remarkable infrastructure of its own. Electronic health records became the system of record for clinical information. Claims systems digitized reimbursement. CRM platforms manage relationships. Data warehouses aggregate information for analysis. Together, these systems transformed how healthcare stores and analyzes information. What healthcare never built was the equivalent infrastructure for communication.

That observation has shaped almost everything we've built at Infinitus. We are building healthcare's communications layer: infrastructure that preserves memory, learns from every interaction, orchestrates what happens next, and gives every participant enough context to move the patient forward without asking them to begin again. That may sound ambitious for something as ordinary as a conversation, but conversations have always carried far more value than healthcare has been able to preserve.

- Patients don't think in organizations.

  They don't see physicians, payors, specialty pharmacies, manufacturers, or financial assistance programs. They see one journey: getting better. Every symptom, appointment, prescription, phone call, and text message is simply another step along the way.

- Healthcare sees something different

  Every organization defines the journey around the moment it becomes responsible for the patient.

- For a physician, that's the first visit.

  For a health plan, enrollment. For a manufacturer, the prescription. For a specialty pharmacy, the referral.

None of those moments is the beginning. They're simply where each organization joins a story already in progress.

That becomes obvious in the conversations themselves. Patients explain their story to a physician, then a nurse, then a specialty pharmacy, then a payor. Every conversation is necessary. Every participant genuinely needs the information. Yet what gets lost usually

isn't the facts. It's the context. Healthcare preserves records, but records aren't memory. An electronic health record knows what medication was prescribed and when a claim was submitted. It rarely knows why a patient is hesitant to begin therapy, when they're easiest to reach, or how weeks of uncertainty have shaped the way they think about their care. Those details rarely fit into structured fields, yet they're often what determines whether the next conversation begins with understanding or with reconstruction.

Healthcare has more information about us than almost any other industry, yet often behaves as though every interaction is the first. Patients repeat themselves. Clinicians retrace decisions. Care teams recreate context that existed only days earlier. The problem is that healthcare has historically treated conversations as transient events. They completed a task and disappeared. In reality, every conversation completes work, generates intelligence, and creates context for the next interaction. Historically, only the work survived. I believe preserving that missing context will become as foundational to healthcare as the electronic health record itself. Patients were never meant to be the continuity layer between healthcare organizations. Technology should be.

Why studio exists

One of the lessons that took us the longest to learn was that customers don't all want the same relationship with technology.

Some want to buy a solution and let us build, operate, and continuously improve it on their behalf. Others want us to build the initial experience but gradually take ownership as their teams become more comfortable. Still others want the platform itself because they see AI becoming a core capability inside their organization. None of these approaches is inherently better than the others. They're simply different points along the same journey.

We started thinking in terms of

- ◯ *"Buy"*
- ◯ *"Buy + Build"*
- ◯ *"Build"*

instead of products versus developer tools. Customers should decide how much they want to own, and they should be free to change that decision over time. One organization may begin by asking us to build everything because they want to move quickly, then gradually take ownership as they develop internal expertise. Another may decide their teams are better served by letting us continue to operate the platform. We think both are valid.

Studio, our AI agent builder and orchestrator, exists to support every point along that spectrum, but not because we believe every customer wants to become a software company. It exists because the people closest to a workflow almost always understand it better than we do. They know the edge cases, the operational realities, the clinical nuances, and the organizational constraints that make their business unique. Our job isn't to replace that expertise. It's to give it the right foundation to build on.

That last point is critical because healthcare isn't an industry where anyone should start with a blank page. Every organization shouldn't have to independently solve clinical escalation, symptom triage, adverse event detection, healthcare reimbursement workflows, quality assurance, human handoffs, conversation memory, or workflow orchestration. They shouldn't have to rebuild integrations with EHRs, CRMs, and payor systems. And they certainly shouldn't have to invent their own approach to governance, compliance, evaluation, or safety for AI in regulated healthcare environments.

Those capabilities are part of the platform.

Years of operating millions of healthcare conversations have been distilled into reusable healthcare building blocks. Clinical escalation pathways. Healthcare-specific orchestration patterns. Our agentic response control (ARC) guardrails and simulation frameworks that continuously evaluate conversations for safety, compliance, and quality. Human review workflows. Conversation memory. Omnichannel communication. Evaluation frameworks. Governance controls. Customers don't begin by solving those foundational problems. They begin with them already in place.

That changes where their time gets spent. Instead of investing months recreating infrastructure that every healthcare organization needs, they can focus on the parts of the experience that are uniquely theirs, including the workflows, clinical programs, patient journeys, and operational innovations that differentiate them. Studio doesn't just make development faster. It lets customers start several years ahead of where they otherwise would.

Some customers will buy solutions. Some will buy and build. Some will build from the beginning. Our responsibility isn't to force a particular adoption model. It's to make sure every customer, regardless of where they start, benefits from the same foundation: healthcare expertise, operational experience, and production-proven infrastructure embedded directly into the platform.

Healthcare's largest untapped dataset

Healthcare generates an extraordinary amount of data. Electronic health records capture diagnoses, medications, laboratory results, and procedures. Claims systems record reimbursement. CRMs track relationships. Data warehouses aggregate everything into reports and dashboards. Nearly every healthcare organization has invested heavily in understanding what happened.

Yet some of healthcare's most valuable information has never lived in any of those systems. It lives inside conversations.

Every day, patients explain why they're hesitant to start therapy. Physicians describe why they chose one treatment over another. Nurses uncover barriers after discharge. Specialty pharmacies discover reimbursement

challenges before they appear in a report. Payors explain denials. Adverse events, product complaints, competitive dynamics, patient sentiment, workflow failures, and emerging opportunities all surface naturally because people talk about them. Healthcare has always generated this intelligence. It simply hasn't had a practical way to capture it. Conversations solved the immediate problem, were summarized in a few notes, and disappeared.

As we automated more conversations, we realized we weren't just helping customers complete work faster. We were creating one of healthcare's richest operational datasets. Every conversation completed a task while revealing how the broader system was performing.

That realization led to Lens. Most organizations already know how many calls they answered or prior authorizations they completed. What they don't know is why. Why are patients abandoning therapy? Why are providers frustrated with a reimbursement policy? Why is one region onboarding more slowly than another? Historically, those questions were answered by manually reviewing 2–3% of conversations weeks after they occurred. Lens analyzes every human and AI conversation in near real time, allowing organizations to detect emerging risks, operational bottlenecks, patient sentiment, and workflow failures as they happen; not weeks later after thousands of patients have experienced the same problem.

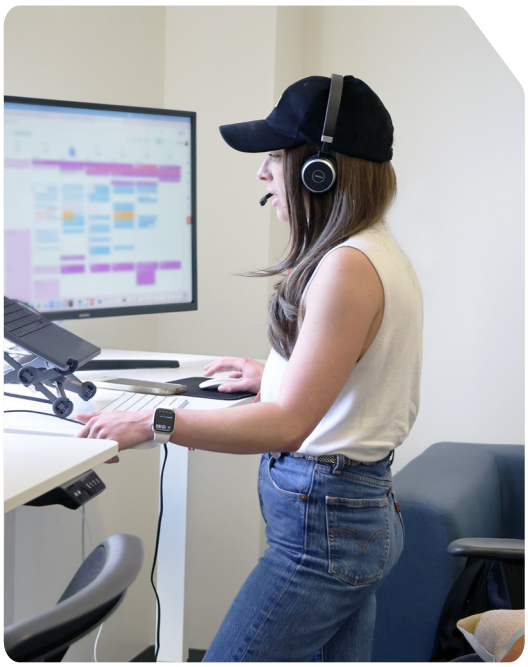

More importantly, the insights don't stop at observation. They become part of a continuous learning loop. Organizations improve workflows, refine programs, adjust policies, and make their AI agents better based on what patients, providers, and care teams are saying every day. Every conversation doesn't just complete today's task. It improves tomorrow's.

I've come to believe this is one of AI's biggest contributions to healthcare. For decades, conversations were ephemeral. Today, they have become one of healthcare's most valuable feedback loops. Healthcare has always been talking. For the first time, it can truly listen and continuously learn.

The technology should experience the same journey the patient does

Imagine you've just been diagnosed with a chronic illness. Your physician recommends a new therapy. Later that day, you receive a text reminding you to schedule an infusion. You call with a question about side effects. A specialty pharmacy emails instructions for getting started. The next morning, your insurance company asks for additional information.

To you, none of this feels like four different interactions. It's one conversation. You're simply trying to get the care you need. Technology should experience the same journey the patient does.

When technology understands the journey instead of the individual interaction, it stops treating every message as a new task. It knows what has already happened, what questions have already been answered, and what is likely to happen next. A phone call isn't just a phone call. It's the next chapter in an ongoing story.

That's why we stopped thinking about channels as products. Voice, SMS, email, chat, patient portals, and whatever comes next are simply different ways of participating in the same conversation. Patients move effortlessly between them. Technology should too.It has rationed attention. Scarce attention inevitably becomes waiting. Waiting for a screening. Waiting for a referral. Waiting for blood work. Waiting for prior authorization. Waiting to begin therapy. Waiting for someone to call back. Healthcare measures those delays operationally—days to therapy, turnaround times, service levels.

More broadly, I think this is where software is heading. For decades, people adapted to software. We memorized passwords, navigated menus, filled out forms, and learned how computers wanted us to behave. Increasingly, software is learning how people behave instead. We speak. We ask questions. We change our minds. We pick up where we left off. We don't think in applications or workflows. We think in goals.

Healthcare should work the same way. When someone asks for help, the technology shouldn't first decide which application owns the request or which channel they're using. It should understand where they are in their journey and respond accordingly. A patient shouldn't have to remind the system what has already happened or explain why they're calling. The system should know a refill is approaching, recognize when someone is at risk of dropping off therapy, anticipate the questions that naturally arise while waiting for treatment, and reach out before uncertainty becomes disengagement. The mechanics of the interaction should become an implementation detail. What matters is that every interaction feels informed, personal, and timely; as though the healthcare system understands the patient rather than simply responding to the latest request.

I don't think the future of healthcare is voice, chat, SMS, or whatever interface comes next. Those will continue to evolve. What shouldn't change is the continuity of the experience.

Patients should feel like they're having one conversation with one healthcare system, even if dozens of organizations, humans, and AI agents are participating behind the scenes.

The best agent isn't always AI

Consider all of the conversations healthcare wishes it could have but simply never does.

A nurse would like to check on every patient after discharge. A specialty pharmacy would like to reach out before a refill is missed instead of after adherence becomes a problem. A physician would spend more time answering questions if the waiting room weren't already full. Care coordinators know exactly which families would benefit from another conversation. Those conversations don't fail because they aren't valuable. They fail because there has never been enough capacity to make them possible.

That's why I've never found the question "Will AI replace people?" particularly interesting. It assumes every task already has a human performing it today. In healthcare, that simply isn't true. There are never enough physicians, nurses, pharmacists, or care coordinators to provide every patient with the experience we all wish they could receive. The constraint isn't compassion. It's capacity.

Sometimes the best agent isn't AI because it's replacing a human. Sometimes the best agent is AI because the alternative isn't a human at all; it's nobody. That's the form of abundance that excites me most. AI doesn't just automate existing work. It makes entirely new kinds of care economically possible by creating attention that never existed before.

That perspective has shaped almost every product decision we've made at Infinitus. We don't believe every conversation belongs in the hands of an AI, just as we don't believe every conversation belongs in the hands of a human. Some interactions require empathy, judgment, and trust. Others are repetitive, structured, and administrative. Automating those conversations isn't about removing

people from healthcare. It's about allowing people to spend more of their time doing the work that only they can do.

It also shapes how we think about the platform itself. We don't optimize for automation. We optimize for the patient. Sometimes the right participant is an AI agent. Sometimes it's a nurse, physician, pharmacist, or care coordinator. The platform's responsibility is to understand the journey, preserve the context, and bring the right participant into the conversation at the right moment. Sometimes that means automation. Sometimes it means escalation. Sometimes it means getting out of the way.

We're working toward a future where patients won't remember whether a person or an AI completed a task. They'll remember whether someone reached out before they fell through the cracks, whether the system understood their situation, and whether they felt supported throughout their journey.

Orchestrating the journey

One of the mistakes people make when they first look at healthcare is assuming the patient journey is a single workflow. It isn't. It's hundreds of workflows owned by different organizations, each optimized for its own responsibilities and incentives.

We came to realize that automating individual conversations was never going to be enough. A perfectly executed benefit verification doesn't help if the next step never happens. A flawless patient outreach call creates little value if the information never reaches the care team. Improving one workflow while ignoring everything around it is like making one relay runner faster while leaving the baton exchange untouched.

That's why the Infinitus platform needed to do more than automate individual tasks, it had to connect and manage the entire workflow.

To us, orchestration isn't about moving boxes around on a workflow diagram. It's about understanding where a patient is in their journey, knowing what should happen next, and ensuring the right participant, human or AI, takes the baton with the full context of everything that has already happened. Sometimes the next step is another AI agent. Sometimes it's a pharmacist reviewing an adverse event. Sometimes it's escalating to a nurse because a patient's symptoms have changed. Sometimes it's simply recognizing that nothing should happen until a physician makes a decision. The platform isn't trying to automate every step. It's trying to coordinate every step.

That's also why we think about workflows differently from most software companies. A workflow isn't a sequence of tasks. It's a sequence of decisions. Every decision

changes what should happen next, who should become involved, what information they need, and how quickly they need it. The more intelligently those decisions can be orchestrated, the more the entire healthcare system begins to feel like a single coordinated team instead of a collection of independent organizations.

Ultimately, patients don't remember individual workflows. As previously mentioned, they remember whether they had to repeat themselves, whether someone followed up when they expected, and whether the next person they spoke with already understood what had happened before. That's the standard we're building toward. Not better workflow automation, but better coordination across the entire patient journey.

Healthcare is different

One of the easiest mistakes to make in AI is assuming healthcare is just another vertical. Build a capable model, connect it to a few systems, add a workflow engine, and point it at a new industry. That approach works surprisingly well elsewhere. It breaks down quickly in healthcare because healthcare isn't just another set of APIs or business processes. It has its own language, clinical workflows, regulatory requirements, and thousands of edge cases that determine whether an interaction remains administrative or suddenly becomes clinical. A refill request becomes a symptom assessment. A benefits inquiry becomes an adverse event report. A scheduling call reveals a patient who needs immediate attention. The technology has to recognize those transitions instantly and know exactly what should happen next.

We learned early that building healthcare AI wasn't primarily about building better models. It was about encoding healthcare itself. Every production deployment taught us something: how clinical escalations should work, how symptom triage should be performed, how adverse events should be recognized, when humans should enter the conversation, how quality should be measured, and how safety and governance should be continuously enforced. Our patent-pending Agentic Response Control (ARC) and evaluation framework grew from that philosophy. Rather than treating safety and compliance as gates before deployment, we built them into the platform so every conversation can be evaluated while it's happening.

That accumulated experience fundamentally changes what customers build. Generic AI platforms begin with an empty canvas. Before solving their own problem, customers first have to teach the platform healthcare—clinical escalation pathways, governance models, evaluation frameworks, human handoffs, and the countless edge cases that only emerge in production. At the same time, their teams are learning how probabilistic AI behaves and how to operate it safely. Everyone is learning at once.

We don't think customers should have to rebuild foundations that already exist. Studio was designed with healthcare primitives built in. Clinical symptom triage based on established frameworks like Schmitt-Thompson, adverse event detection, conversation memory, human handoffs, orchestration, healthcare integrations, governance, and ARC guardrails are available from the start. Every agent can then be evaluated against a rubric developed by our clinical team of physicians and nurses and stress-tested through hundreds or thousands of simulated conversations spanning different patient personas, clinical scenarios, and edge cases before it ever reaches a real patient.

That changes where customers spend their time. Instead of teaching AI how healthcare works, they can focus on the workflows, patient journeys, and experiences that make their organization unique.

The platform absorbs the common complexity so customers can focus on the complexity that actually creates value for patients.

The Infinitus platform

If you've read this far, you've probably noticed that I rarely describe Infinitus as an AI company. That's intentional. AI is one of the technologies we use, but it has never been the organizing principle behind the platform. The organizing principle is much simpler: every patient journey is a series of conversations, and every conversation should leave the healthcare system better prepared for the next one.

At first glance, a conversation seems like a simple exchange of information. In reality, every conversation gives the platform a new responsibility. It has to communicate with the right person in the right way. It has to remember what was learned so the next interaction doesn't begin from zero. It has to learn from millions of similar conversations to identify operational bottlenecks, patient needs, and opportunities to improve care. It has to decide what should happen next and who should be involved. It has to do all of that safely, with governance, compliance, quality evaluation, and clinical oversight built into the foundation rather than added afterward. Finally, it has to give every healthcare organization the ability to shape these capabilities around its own workflows instead of forcing them into ours.

Looking back, what we were really building wasn't a collection of healthcare workflows. Every new use case, including benefit verification, prior authorization, patient onboarding, refill reminders, symptom triage,

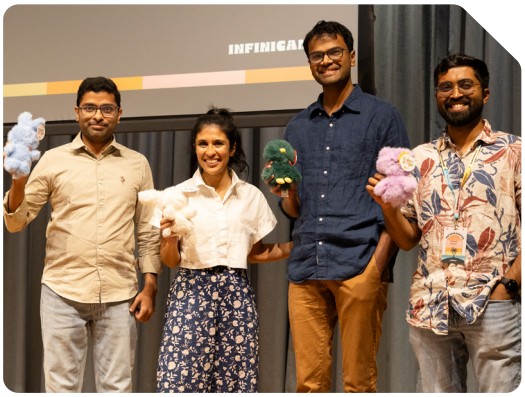

clinical escalations, and countless others, relied on the same underlying capabilities. The platform had to understand what someone was asking, apply the right clinical and business rules, preserve context, coordinate handoffs, document outcomes, know when to escalate to a human, and learn from every interaction. Individual workflows came and went. Those capabilities kept getting reused. That's when we realized we weren't building solutions for individual problems. We were building the foundation they all shared.

> We didn't discover those capabilities by sitting in a conference room drawing a platform diagram. We discovered them because customers kept trusting us with the problem next door.

Every enterprise relationship begins with a single workflow. Someone inside the organization spends political capital, convinces skeptical colleagues, and puts their own credibility behind a company they've only known for a few months. If that first deployment succeeds, the conversation changes. Customers stop asking whether you can solve the original problem and start asking whether you can solve the adjacent one.

- ◯ *Could you automate prior authorization?*
- ◯ *Could you support inbound patient calls?*
- ◯ *Could we build our own agents?*

Looking back, it appears as though the platform followed a carefully designed roadmap. Living through it, it felt much simpler. Every new problem revealed another capability we needed to build, and every capability became reusable across the next set of workflows. We didn't invent the platform. Our customers revealed it to us, one adjacent problem at a time.

Communication is where everything begins. Patients don't think in channels, and they shouldn't have to. Whatever channel an interaction begins on shouldn't change the experience because what matters is the conversation, not the channel. That conversation then becomes memory. Every interaction should preserve the context that allows the next participant, whether human or AI, to begin with understanding instead of rediscovery. At the same time, every conversation becomes intelligence. Across millions of interactions, patterns emerge that reveal where patients struggle, where workflows break down, and where organizations can improve. That's the role Lens plays within the platform.

Conversations, however, are only meaningful if they move the patient forward. A benefit verification leads to a prior authorization. Prior authorization leads to enrollment. Enrollment leads to scheduling. Treatment leads to adherence. Every step creates another decision, another handoff, and another opportunity for the patient to fall through the cracks. That's why orchestration became such a central part of the platform. The goal isn't simply to automate tasks. It's to understand where the patient is in their journey, preserve the context surrounding every decision, and coordinate the right next participant. Sometimes that's another AI agent. Sometimes it's a pharmacist reviewing a clinical escalation. Sometimes it's a physician, nurse, or care coordinator. The platform

shouldn't be biased toward automation. It should be biased toward the patient.

None of this would matter if customers had to rebuild these capabilities from scratch every time they wanted to create something new. That's why Studio exists. We don't believe healthcare organizations should spend months reinventing clinical escalation pathways, symptom triage, adverse event detection, governance models, quality evaluation, ARC guardrails, conversation memory, or human handoffs. Those are foundational healthcare capabilities that belong inside the platform itself. Customers should be free to build what makes their organization unique while standing on infrastructure that's already learned from millions of healthcare conversations. Some customers want us to build everything for them. Others want to extend what we've built. Others want to build and manage their own experiences from the beginning. We think all three approaches are valid, which is why you'll recall we've always described the platform as supporting **"Buy," "Buy + Build,"** or **"Build."**

When I step back, I don't see a collection of products anymore. I see a platform with a clear set of responsibilities. It has to communicate. It has to remember. It has to learn. It has to orchestrate. It has to govern. And it has to empower customers to shape all of those capabilities around the realities of their own organizations. Every product we've built exists because one of those responsibilities couldn't be ignored.

That's also why I'm confident the platform will outlive today's generation of AI models. The models will change. The interfaces will change. The underlying technology will continue to evolve. But healthcare will always need to communicate, preserve context, coordinate across organizations, learn from every interaction, and do all of it safely. Those aren't product features. They're foundational capabilities.

Ultimately, that's how I hope people think about Infinitus. Not as a voice company, an AI company, or even a software company. I hope we're remembered as the company that built the communications infrastructure healthcare never knew it was missing, helping the entire system behave a little more like the single journey every patient assumed it already was.

ACT III

Why now and why us

Why now?

Great companies are often built when three things become true at the same time: the technology is finally capable, the market is finally ready, and the buyer finally has both the urgency and authority to act.

I think healthcare has reached that moment.

1. **The technology is ready.**
   AI can now understand natural conversation, reason across complex workflows, communicate across voice, text, email, and chat, and continuously improve through experience. Ten years ago, building a healthcare communications layer simply wasn't practical. Today it is.

2. **The industry is ready.**
   AI is no longer an experiment. Nearly every major health system, payor, specialty pharmacy, and life sciences company has concluded that AI will fundamentally reshape healthcare. The question is no longer whether to adopt it. It's how fast.

3. **The buyer has changed.**
   The demand is no longer coming primarily from innovation teams or centralized IT organizations. It's coming from business leaders who own P&Ls and are accountable for access, adherence, patient experience, and administrative cost. They don't need another AI strategy. They need production results.

That has created one of the defining debates in healthcare today: buy or build?

Almost every large organization wants AI to become a core competency, and they should. But production-grade healthcare AI requires much more than access to a frontier model. It requires clinical safety, governance, evaluation, healthcare-specific workflows, integrations, orchestration, and millions of conversations worth of learning. Internal teams are being asked to deliver all of that while the business is asking for results this quarter.

That's why I don't think the future belongs to companies that only sell finished applications or only sell developer tools. Healthcare needs both. Some organizations want to buy. Others want to build. Most will do both. The vendors that win will provide the healthcare primitives that let customers move quickly without rebuilding decades of healthcare complexity.

The question isn't whether healthcare will adopt AI. It's who the industry will trust to build upon.

Why us?

Every major technology platform eventually creates a new generation of vertical companies. Cloud computing didn't eliminate healthcare software. It made healthcare software more valuable. Salesforce didn't replace industry-specific CRMs. It became the platform they were built on. Foundation models will follow the same pattern.

The large horizontal AI companies will build extraordinary infrastructure. They should. They will continue pushing the frontier of models, compute, developer tools, and general-purpose agents. We use many of those technologies ourselves. But healthcare isn't a generic domain that happens to use AI. It's a domain where the workflows, regulations, clinical decisions, and operational realities are the product.

A technology platform catering to healthcare must understand why a prior authorization stalls, how a specialty pharmacy thinks about adherence, what constitutes a reportable adverse event, when symptoms require escalation, how physicians, payors, manufacturers, and care teams coordinate, and where patients actually fall through the cracks. Those aren't capabilities you get from a frontier model. They are learned – one workflow, one deployment, and one customer at a time.

We've spent over seven years accumulating exactly that knowledge. Along the way, we discovered that every new deployment required many of the same underlying capabilities: clinical triage, safety escalations, anti-kickback compliance, adverse event detection, conversation memory, human handoffs, governance, quality evaluation, simulations, healthcare integrations, and orchestration. We built these primitives into the Infinitus platform itself. Today our customers don't begin by teaching an AI how healthcare works. They begin with healthcare already built in.

That's the difference between building on a horizontal AI platform and building on a healthcare platform. A horizontal platform gives you intelligence (raw material). A vertical platform gives you judgment (outcome). It gives you workflows that have already been proven, governance that has already been built, clinical safety that has already been evaluated, and infrastructure that has already learned from millions of healthcare conversations. Instead of spending months assembling the foundations, customers can focus on the workflows and experiences that make their organizations unique.

I don't think we'll win because we have the best model. The best model will change every year. I think we'll win because we understand healthcare more deeply every year. Models are becoming commodities. Healthcare expertise compounds.

ACT IV

Building with intention

Choosing what to build

Every company eventually discovers that ideas aren't the scarce resource. Time is.

Customers ask for another feature. Employees see another market to enter. Investors suggest adjacent opportunities. Every new technology creates another set of possibilities. Especially today, AI has made the number of things we could build almost limitless. The hard part isn't finding ideas. It's deciding which ones deserve years of your life.

I've found myself coming back to the same questions over and over again, although I didn't recognize them as questions at first.

The first is whether the problem actually matters. I don't mean whether it sounds ambitious or whether it's attached to a large market. I mean something much simpler. Does solving this problem genuinely improve the patient journey? Does it help someone receive care sooner? Reduce uncertainty? Remove friction that never should have existed? If I can't draw a direct line from the thing we're building to a better experience for a patient, it's probably not something I want us to spend years building.

That sounds obvious, but it's surprisingly practical. Every roadmap is full of worthwhile ideas. Purpose isn't what tells you what to build. Purpose tells you what to ignore.

The second question is more uncomfortable because it forces honesty. Are we actually the right team to solve this problem? Companies don't win because they pick the biggest opportunity. They win because they become uniquely good at solving it. Choosing a difficult problem is really choosing the company you're going to become. It determines who you hire, what expertise you accumulate, and what you're willing to spend years learning before it compounds into an advantage.

The final question is the one that keeps companies alive. Will customers continue asking us to solve this problem? I've never believed enduring companies are built by raising money forever. The strongest companies eventually earn the right to fund their own future because customers value what they've built enough to keep expanding

it. One of my favorite moments in any customer relationship isn't when they buy the first product. It's when they ask us to solve the second problem. That's usually the moment you've stopped being a vendor and started becoming a partner.

Looking back, I don't think these are really three separate questions. They're different ways of asking the same thing. Is this a problem worth dedicating a meaningful part of your life to? If the answer is yes, then it's worth becoming the best team in the world at solving it. If you succeed, customers will usually tell you what to build next.

Infinitus is the communications platform for healthcare. AI is how we build it today, but it's not why we exist. We exist to improve the patient journey. The technology will change. The patient journey won't.

Customer obsession

I've never believed that being customer obsessed means giving customers everything they ask for. In fact, that's one of the fastest ways to build a mediocre product.

Customers are incredibly good at describing friction. They know exactly where a workflow feels slow, repetitive, or confusing. They know where they lose time. They know where people become frustrated. What they aren't always able to do is describe the root cause. That's our job.

I've found that the most valuable part of a customer conversation usually begins after the feature request. Someone asks for a new capability, but as the discussion continues you realize the feature wasn't really the point. It was simply the easiest way they knew how to describe a deeper problem. That's why I like asking "why?" one more time than feels necessary. Why does this matter? Why is this the process today? Why didn't the last solution work? Those questions rarely lead directly to the product we build. They help us understand the problem well enough to imagine solutions the customer never thought to ask for.

Customers rarely tell you what product to build. They tell you where they're struggling. Our job is to understand that struggle well enough to imagine solutions they couldn't reasonably ask for because, until recently, the technology simply didn't exist. The conversation becomes less about collecting requirements and more about expanding what's possible.

> Customer obsession also means listening for what isn't being said.

Sometimes it's the hesitation before someone answers. Sometimes it's the excitement when they describe a future workflow instead of their current one. Sometimes it's noticing that three different customers are describing the same frustration using completely different language. Those are often the moments where the next product idea comes from.

It also means knowing when to disagree. Customers understand their problems better than anyone. That doesn't always mean they know the best solution. If we immediately say yes to every request, we learn very little. Sometimes the most valuable thing we can do is push back, not to win an argument but to

better understand how strongly they believe in the problem they're describing. If they immediately move on, it probably wasn't that important. If they passionately explain why it matters, we've learned something we couldn't have learned by simply agreeing.

Over time, I've realized that customer obsession isn't really about customers. It's about curiosity. It's believing that every conversation contains more information than was actually spoken. It's understanding that our job isn't to build exactly what customers ask for. It's to understand the future they're trying to describe before they have the words to describe it themselves.

Judgment and taste

Building a company is ultimately an exercise in making decisions with incomplete information. Markets change. Customers surprise you. Competitors appear overnight. If you wait until every answer is obvious, you've already fallen behind.

That's why I've never thought great leaders are the ones who make perfect decisions. I think they're the ones who learn the fastest. They form hypotheses, test them against reality, update quickly, and repeat. Looking back, it appears they had remarkable instincts. More often, they simply iterated faster than everyone else.

That requires holding two ideas at the same time. I like forming strong opinions because they create clarity. I like holding them weakly because new evidence should always be allowed to win. A mentor once told me that if someone changes your mind in five minutes about something you've supposedly spent months thinking about, you probably hadn't thought deeply enough about it in the first place. Strong opinions force rigor. Weakly held opinions force humility. Good judgment requires both.

Judgment, however, isn't enough. Startups rarely fail because they have too few ideas.

They fail because they chase too many. Every customer asks for another feature. Every new technology creates another opportunity. Every roadmap becomes longer than the company has time to build. That's where taste becomes essential.

People often associate taste with aesthetics. I think it's much broader than that. Taste is knowing what belongs and, more importantly, what doesn't. It's recognizing the few ideas that reinforce the identity of the company and having the discipline to say no to the hundreds that don't.

Focus isn't simply resource allocation. It's preserving coherence.

Over time I've come to think that judgment and taste reinforce one another. Judgment helps you change course when reality changes. Taste prevents you from changing course every time the wind blows. One keeps the company adaptable. The other keeps it recognizable.

The companies I admire most have both. They learn quickly without becoming directionless. Their products evolve, but they still feel like they came from the same mind. Their strategy changes, but their philosophy doesn't.

Momentum

People often ask how I know whether we're winning. The obvious answers are revenue, customers, product launches, or hiring. While those metrics matter, they usually lag reality. By the time they show up on a dashboard, something has already changed inside the company.

I believe momentum is something you feel before you measure. The role of metrics isn't to replace that intuition. It's to validate it early enough that you can invest with conviction.

You feel it when customers start asking for more than they originally bought. You feel it when candidates walk into interviews already convinced they want to join. You feel it when people who never had time for you suddenly make time because they want to be part of what you're building. You feel it when a product demo stops being a presentation

and starts becoming a brainstorming session because customers are already imagining what's next.

The interesting thing about momentum is that it feeds itself. Success creates belief. Belief creates energy. Energy creates better

execution. Better execution creates more success. At some point it becomes difficult to separate cause from effect because everything starts reinforcing everything else.

I've also never believed that burnout comes from working hard. Some of the hardest periods of my career have also been the most enjoyable because we could see progress every week. Long hours are surprisingly energizing when they're producing visible results. What's exhausting is working just as hard while feeling like nothing is changing.

When effort stops producing progress, people begin questioning the journey. When effort produces momentum, they want to push even harder.

Momentum also raises the standard. Customers expect more because they know you're capable of more. Teams expect more from each other because they've seen what's possible. Success doesn't make people complacent. It creates a responsibility to keep improving.

Leadership

One of the biggest surprises I've had as a founder is how many responsibilities I assumed would eventually become someone else's. Some did, but many never have.

As a company grows, decisions become more distributed and teams become more specialized. That's exactly what should happen. But some responsibilities never leave the founder's desk because they define the company.

Product taste is one. Culture is another. Hiring is another. For as long as I can, I want a personal relationship with every employee because culture spreads through people, not values written on a wall.

The same is true for customers. The fastest way to lose product taste is to become insulated from the people using what you build.

I've also never believed in pure people managers.

Everyone at Infinitus is expected to be a player-coach. Leaders should still build, sell, recruit, design, and solve problems. The moment your job becomes telling other people what to do instead of making the work better, you've drifted too far from the work itself.

As companies grow, every function naturally optimizes for itself. Sales wants one thing. Engineering wants another. Marketing sees a different opportunity. Leadership is making sure everyone is still playing the same game.

Culture

I've come to believe that culture is what people consistently do when nobody tells them what to do.

One of the qualities I value most is agency. Great companies aren't built by people waiting for permission. They're built by people who see a problem, assume ownership, and move. They don't ask, "Whose job is this?" They ask, "How do we solve it?" Agency creates momentum because problems don't sit around waiting for the perfect owner.

Agency without urgency, however, doesn't accomplish much. Healthcare is full of patients who are waiting. Waiting for answers. Waiting for approvals. Waiting to begin therapy. Every unnecessary day matters. That creates a bias toward action inside our company. We move quickly, not because speed is a virtue on its own, but because the work we do affects people who don't have the luxury of waiting.

Just as important is fearlessness. I don't mean taking reckless risks. I mean being willing to tell a customer something they don't want to hear, challenge conventional wisdom, walk away from opportunities that don't fit, or tackle problems everyone else assumes are impossible. Fearless teams don't confuse confidence with certainty. They simply refuse to let fear become the reason they make a decision. Most breakthroughs look unreasonable before they look obvious.

We also care deeply about humility. The best teams I've worked with care more about finding the right answer than proving they already have it. Strong opinions make companies better. Humility makes those opinions change when reality says they should. That's why we talk so much about

the importance of the "disagree and commit" concept. Healthy teams should debate, challenge ideas, and push back when they think something is wrong. But once a decision is made, everyone owns it. There is no value in winning the debate if the company loses the execution.

Finally, we try to be optimistic realists. You have to believe healthcare can become dramatically better, otherwise there is no reason to spend years trying to change it. But optimism without realism becomes wishful thinking. Realism without optimism becomes complacency. The goal is to hold both at the same time.

I've found that culture isn't really tested when things are going well. It's tested when the product breaks, a customer is unhappy, or a deadline slips. Those are the moments when agency matters. Urgency matters. Fearlessness matters. Humility matters. Optimism matters. That's when you discover whether your culture is real or just something written on a wall.

Endurance

I've never been interested in building a company just to make it to the next funding round. My goal has always been to build something that still matters 30 years from now. Operating on that kind of timeline shifts almost every decision you make. You stop obsessing over what customers will buy this quarter and start asking what they will fundamentally need a decade in the future. You stop chasing every passing tech trend and start identifying the core problems that will persist no matter how technology evolves.

One of the hardest lessons I've learned along the way is that simply being right isn't enough. You also have to survive long enough for the market to agree with you. Looking back, we definitely pursued certain ideas before our customers were truly ready for them. We walked away from some too early, and probably held onto others for too long. Yet, while the timing fluctuated, the underlying direction rarely changed. We kept coming back to the exact same belief: that healthcare would eventually have to become more connected, automated, and intelligent. The timeline shifted, but our conviction didn't. There were moments when it would have been easier to pivot toward whatever the market was excited about at the time. We chose instead to keep investing in the same fundamental problem because we believed the market would eventually meet us there.

It is also critical not to mistake external noise for progress. Customers, investors, and competitors will always overwhelm you with their opinions. While those signals are worth tracking, they shouldn't dictate your compass. Conviction isn't a stubborn refusal to change your mind. It's knowing exactly what should never change.

For us, the products have evolved, and the technology certainly has, but the foundational belief that every patient deserves a simpler, faster journey remains completely non-negotiable. That's why we've always tried to build our company around problems instead of specific technologies. Technologies evolve. Patient needs don't. The companies that endure over decades aren't the ones that predict the future with perfect accuracy;

they are the ones that stay deeply committed to solving the structural problems that will still matter whenever that future finally arrives.


Infinitus