

AI-Powered Patient Access:


An Evaluation Framework for Healthcare Organizations

What's Inside

- Page 3 **Introduction**
- Page 5 **The Access Problem Healthcare Leaders Are Trying to Solve**
- Page 7 **The Growing Noise Around Healthcare AI**
- Page 8 **Why Healthcare Scheduling Is More Complex Than It Appears**
- Page 10 **What Healthcare Organizations Should Evaluate in an AI-Powered Patient Access Platform**
- Page 12 **Security, Governance, and Operational Guardrails**
- Page 14 **Data Ownership, Infrastructure, and Integration Considerations**
- Page 15 **The Evaluation Framework: A Buyer's Scorecard**
- Page 16 **The Case for Disciplined Evaluation**

Every organization evaluating AI-powered patient access is really answering two questions at the same time: what an AI solution can do, and what it actually needs to do. The first question is easy to ask. The second one is easy to skip.

That second question is the one that makes the difference. And it is getting easier to skip, because the market keeps getting louder. According to [Fierce Healthcare](#), half of healthcare organizations now run three or more AI solutions, and every one of them demoed well and promised to ease the same familiar pressures. When a solution clears the evaluation on capability alone, an organization can end up with a genuinely capable tool that might not be a match to the problem it needed to solve, the workflows it had to fit, or the oversight the team expected once it went live. What looks like a technology decision is often really a business decision that's worth pausing to consider first. And the deciding factor usually is not the technology itself. It's the evaluation.



What **Evaluating AI-Powered Patient Access Involves**

Evaluating an AI-powered patient access solution means looking at how well it handles the full complexity of patient access, not just a simple interaction. That includes scheduling logic and rules depth, integration with EHR and practice management systems, coordination across phone, digital, and self-service channels, security and governance controls, operational reporting and auditability, and the vendor's ability to support the solution as your operations grow. The strongest evaluations start by defining your own access challenges first, then measuring each solution against it.

“

AI is an answer. Make sure you know the question. If you have areas of pain, have you measured them? What have you already tried? What work is already in progress to address that pain? Know what you're solving for before you decide what solves it.



David Dyke


Chief Product Officer
Relatient

The Access Problem **Healthcare Leaders Are Trying to Solve**

Before any solution enters the conversation, it helps to be honest about the reality of patient access today. Providers are stretched thin, and burnout is real. Front-desk and administrative roles are hard to fill and harder to keep staffed. Contact centers strain under call volume that spikes without warning. Patients call after hours and expect someone, or something, to help. And underneath all of it sits scheduling, which is far more complicated than it looks from the outside.

It is tempting to treat these as one big challenge called “access.” They’re not. They’re distinct challenges, and they do not share a single solution. A tool that eases one may do nothing for the other, which is why the first question is not which solution to buy. It’s which challenge you actually have. Is it volume, staffing, peak times, after-hours coverage, competency, or the pace of change itself?

Most organizations have more than one challenge, and the point is not to pick a single label but to see them clearly. That clarity is what makes the difference, because it separates the challenges that need more resources from the ones that need better process. Often what looks like a resource constraint is really workload, with routine manual tasks consuming hours that staff can spend elsewhere. Naming the problem precisely is what turns a vague sense of strain into something you can act on, and something you can measure, from throughput to workload and consistency.



Which problem do you need to solve?

1

Is it a volume issue, where more inbound calls come in that your team can answer?

2

Do you have the staff or roles but struggle to fill or keep them?

3

Are patients reaching out after hours with no one available to help?

4

Are patients getting booked accurately, with the right provider at the right location?

5

Are provider schedules unbalanced, with new providers sitting empty while established ones are overbooked?

The Growing Noise Around Healthcare AI

For years, healthcare was the industry that waited, mired in regulation and a full generation behind every major technology wave. AI has flipped that pattern. According to Menlo Ventures' State of AI in Healthcare report, healthcare organizations are now adopting AI solutions 2.2 times faster than the broader economy, shedding a long-held reputation as a technology laggard. The market has filled with vendors just as quickly. That energy is real, and it signals genuine promise. It also means the hardest part is no longer finding a solution. It's telling strong ones apart.

With so many options arriving at once, many look alike on the surface and demo well, which makes it easy to move quickly, and just as easy to move before defining what your organization needs. A fast-moving market rewards buyers who slow down long enough to know what they are solving for.

The advantage grows with the solution. The more a system can do, the more thoughtfully your organization can put it to work, deciding what capabilities to use, what to hold back, and who owns the result. That intentionality is what turns a capable tool into a well-run one, and it is what separates organizations that get lasting value from those that deploy more than they planned in places they never fully evaluated. The strongest adopters treat a long feature list as an invitation to be deliberate. This is simply the discipline that makes early investment pay off. When decisions move quickly, clear evaluation is what protects the value of the choice over time. The organizations that evaluate carefully are the ones that see durable results, because they matched the technology to a problem they understood first. That is the difference between buying with the market and buying with intention.



“Long dismissed as a digital laggard, healthcare is now setting the pace for enterprise AI adoption.”

*2025 State of AI in Healthcare Report
(via Menlo Ventures)*


Why Healthcare Scheduling Is More Complex Than It Appears

Of everything an AI-powered patient access platform is asked to do, scheduling is the hardest to get right and the easiest for a vendor to make look simple. In a demo, booking is a clean moment. A patient asks, the system finds an available slot, and the appointment appears. But that smooth surface hides the part that matters, which is everything the system has to know before it offers that slot.

Real scheduling runs on provider preferences, and there are far more of them than most organizations realize. One physician might not take new patients on Friday afternoons. Another sees post-operative follow-ups only in the mornings. A third wants new Medicare patients spaced out. These rules and preferences live in the head of schedulers who learned them over the years, so a platform that cannot capture them does not simplify the work. It just moves the errors somewhere new.

The complexity compounds from there. Appointment types must route to the right provider, location, and visit duration. Referrals carry their own authorizations and criteria. Multi-location and time zone coordination means knowing who works where and when. Sequential visits have to happen in order, like imaging before a consult. And bulk scheduling has to do all of it at once when a provider's day shifts, without breaking a single rule underneath.

This is why “can you create new appointments?” is a flawed question for vendors. The answer is always yes. What separates a capable platform from a convincing demo is whether it can create the right appointment, for the right patient, with the right provider, under all the rules an organization runs on. The scenarios that follow show where that complexity lives, what it risks when mishandled, and what it delivers when done right.



Where **Scheduling Complexity Lives**

Scheduling Capability

- 🕒 Provider-specific rules
- 🕒 Appointment-type routing
- 🕒 New patient intake
- 🕒 Referral workflows
- 🕒 Multi-location coordination
- 🕒 Sequential scheduling
- 🕒 Waitlist backfill backfill
- 🕒 Bulk scheduling

Risks if Mishandled

- ⚠️ A mismatched booking may not surface until the visit, leading to a wasted appointment and a reschedule for the patient
- ⚠️ Incorrect visit length or provider, creating gaps and overruns
- ⚠️ Incomplete details stall booking and create bottlenecks at the point-of-care
- ⚠️ Referred patients wait on manual outreach and fall through
- ⚠️ Patients book at the wrong location, causing no-shows
- ⚠️ Dependent visits booked out of order, delaying care
- ⚠️ Canceled slots go unfilled, leaving providers idle
- ⚠️ Schedule changes trigger manual rebooking

Rewards When Done Right

- ⭐ Preferences are respected automatically, protecting provider time
- ⭐ The right visit reaches the right flow, keeping utilization high
- ⭐ Information is captured cleanly so new patients book without friction
- ⭐ Referred patients self-schedule directly, closing referral leakage
- ⭐ Availability reflects which providers work at which locations, booking care where it happens
- ⭐ Linked visits are sequenced correctly, keeping pathways on track
- ⭐ Open slots fill automatically, recovering utilization
- ⭐ Large changes are handled at once, honoring every rule

What Healthcare Organizations **Should Evaluate in an AI-Powered Patient Access Platform**

Once you know the problem you are solving, evaluation becomes a matter of looking in the right places. Most would judge platforms on how naturally they hold a conversation, because that is the part they can see in a demo. But the conversation is the surface. What determines whether a platform delivers over the years, across locations, under real conditions, is a set of capabilities that rarely make it into the pitch. Here are six criteria that matter the most:

1 Workflow Orchestration

The question is not whether the platform can complete a task, but whether it can coordinate the whole journey, routing each request to the right resource, moving work between automation and staff, and handling what comes after the call as reliably as the call itself. A platform that automates one step in isolation leaves the seams for your team to manage.

2 Intelligent Scheduling Depth

As the previous section made clear, booking is where complexity hides. Evaluate whether the platform can hold the real rules your organization runs on, provider preferences, appointment types, triage questions, referrals, sequencing, and multi-location logic, rather than a simplified version that looks right in a demo and breaks in production.

3 Real-Time Integration and Data Synchronization

A platform is only as good as its connection to your source of truth. Evaluate whether it reads and writes to your EHR and practice management systems in real time, so availability, patient records, and appointments stay accurate the moment anything changes. Anything less introduces the very errors automation is supposed to remove.

4 Governance and Security Guardrails

Evaluate what the platform will and will not do on its own, how it escalates, and what controls you have over its behavior. This is covered in depth in the sections ahead, but it belongs on the core list because a platform without guardrails is a liability no capability can offset.

5 Operational Oversight and Reporting

Look at whether you can see what the platform is doing, audit its interactions, and measure its performance the way you would manage a team. Visibility is not a reporting feature bolted on at the end. It is how you keep an automated system accountable.

6 Deployment and Scalability

See how the platform is implemented, who owns its configuration, how updates are handled, and whether it holds up as volume grows and rules change. A platform that works at one location under today's conditions is not the same as one that scales across your organization and keeps pace with it.

These six are easier to satisfy when access is unified rather than assembled. The strongest platforms treat AI-driven channels, digital, and the contact center as one experience, sharing the same rules, data, and view of the patient, so a request that starts on the phone and continues online never loses its thread. Unifying access is also what makes the humans-plus-AI balance possible, because the most effective deployments pair automation with people deliberately, letting AI carry the volume it handles well while keeping humans on the interactions that matter most. Some organizations decide the first contact with a new patient should always be a person. Payments draw a similar line: confirming a copay already on file suits automation, while collecting a new payment method through an AI conversation sits at the far edge of patient comfort. A capable platform supports those choices rather than overriding them.



Security, Governance, and Operational Guardrails

Every other criterion in this framework measures what a platform does. This one measures what it is, and in healthcare, what a system is, secure, governed, accountable, determines whether anything it does can be trusted. A useful principle organizes the rest: AI-powered workflows should be managed with the same rigor as a human employee. No organization would place a new hire on the phones with no supervision and no record of their work, and an AI agent warrants the same expectation. That expectation takes a concrete form:

Auditability

You can examine transcripts, review recordings, and investigate the calls that did not go well.

Defined boundaries

You set which tasks the agent performs, which it must escalate, and where a human assumes control.

A reliable escape hatch

Any patient has a dependable, obvious way to exit the automated experience and reach a person.

Data movement is where organizations most often expose themselves, and it deserves the same scrutiny. Third-party risk now dominates: in the first half of 2026, business associates were involved in 43% of reported healthcare data breaches, according to the [HIPAA Journal](#). Several of the more damaging incidents in healthcare AI originated not in a breach of the vendor itself, but in data quietly reaching a third party the patient never knew existed. The question in evaluation is therefore not simply what a vendor does with patient data, but precisely who handles it, and under what agreement.



Certifications enter here, and they deserve careful interpretation. HIPAA, HITRUST, and SOC 2 are evidence, not proof, a distinction worth stating plainly, because the two are easily conflated. A certification establishes that a vendor has satisfied a recognized standard, which is a floor to be cleared. Treating it as a binary, present or absent, discards the more important information, namely how data is handled, where it resides, who may access it, and whether those practices hold under the specific condition your organization imposes. A certified vendor offering shallow answers warrants more concern than a transparent one still working toward certification. The certification opens the inquiry. It does not close it.

“

Trust in a healthcare technology platform is anchored in the fundamentals of confidentiality, integrity, and availability, and healthcare organizations should expect to see those principles not just asserted but demonstrated in everything the platform does with patient data, every day.



David Edwards
Chief Technology Officer
Relatient

Data Ownership, Infrastructure, and Integration Considerations

Some of the most important questions in an evaluation are the ones buyers rarely think to ask. Where does patient interaction data live? Who owns it, the organization, or the vendor? And is it used to train the vendor's models? These determine whether your patients' conversations remain yours or quietly become an asset on someone else's balance sheet.

These are worth raising directly with any vendor. It can help confirm where data resides and how that infrastructure is controlled, to understand how ownership is defined in the agreement, and to ask whether data is used for model training and, if so, on what terms and with what options to opt out. Vendors take different approaches to each, so the goal is clarity on how a given platform handles them rather than a single right answer.


Integration is another area worth examining closely. Some tools connect to existing systems at a surface level, while others integrate more deeply. It is helpful to understand how a platform exchanges data with your EHR and practice management systems, how current that data stays as things change, and where information passes between tools, since those handoffs can affect accuracy. The right depth depends on your environment and what you need the platform to do.

A few further questions help distinguish a proven platform from an earlier-stage one. Who can access patient data, and how is that access controlled and logged? How does the infrastructure perform as volume grows, beyond pilot scale? These rarely appear in a feature list, yet they shape how a platform performs in practice, which is why integration architecture is worth treating as a core part of the evaluation rather than a detail settled afterward.



The Evaluation Framework: **Vendor Evaluation Scorecard**

*How to use: Set a priority for each item, then score how well each vendor meets it.
Total the scores to compare vendors side by side.*



Evaluation Criteria	Priority	Vendor A	Vendor B	Vendor C
Healthcare-Specific AI Capability				
Trained on healthcare patient conversations, not generic call center data				
Understands specialty-specific workflows				
Escalates to a human agent when the call requires their judgment				
Handles varied speech patterns, accents, and pacing				

Evaluation Criteria	Priority	Vendor A	Vendor B	Vendor C
Appointment and Access Workflows				
Manages the full appointment cycle (book, reschedule, cancel) without staff involvement				
Answers common questions such as directions and clinic hours, and transfers to a human scheduler when needed				
Books appointments for both new and existing patients				
Captures and routes clinical requests such as nurse callbacks, without staff involvement				
Manages waitlists and fills open schedule gaps automatically				
EHR and Systems Integration				
Native integration with your EHR (athenahealth, Veradigm, Epic, Oracle, and others)				
Writes back to the EHR in real time with no batch sync or manual reconciliation				
Creates cases or tasks directly in the EHR				
Identifies the caller and brings up relevant patient context				
Compliance and Security				
HIPAA business associate agreement (BAA) in place				
SOC 2 compliant				
HITRUST certified				
PHI encrypted in transit and at rest				
Configurable boundaries that limit the AI to approved workflows				

Evaluation Criteria	Priority	Vendor A	Vendor B	Vendor C
Performance and Analytics				
Measures staff time saved through automated call handling				
Reports reduction in call abandonment rate				
Tracks containment rate (calls resolved without staff)				
Captures patient satisfaction scores after the call				
Logs the reason and outcome for every interaction				
Provides a full audit log for every call, including transcript and audio				

Evaluation Criteria	Priority	Vendor A	Vendor B	Vendor C
Implementation and Partnership				
Proven go-live timeline in specialty healthcare settings				
Dedicated customer success support with patient access expertise				
Staff training and change management support included				
Documented SLAs and uptime guarantee				
Advisory support available for workflow design and optimization				

Total Score

The Case for Disciplined Evaluation

This guide sets out to make one shift possible moving from evaluating AI on how well it demonstrates to evaluating it on how well it will actually work. That means starting with your own access challenge rather than a vendor’s feature list, looking past the smooth surface of a demo to the scheduling logic, integration depth, governance underneath, and treating security and visibility as the floor rather than a bonus. It means recognizing that a capable platform and the right platform are not the same thing, and that the difference is found through disciplined evaluation, not impression.

What Disciplined Evaluation Looks Like:



Starting with understanding your own access challenge. AI is an answer and the first job is knowing the question.



Judging depth over demos, testing the scheduling logic, and governance that a polished pitch does not show.



Demanding visibility and control, managing an AI agent the way you would manage staff, with audit, boundaries, and a reliable escape hatch.



Holding every vendor to the same standard, using the scorecard to compare platforms side by side rather than by impression.



“

The healthcare organizations that get the most from AI-powered patient access solutions are the ones that take a disciplined, long-term approach rather than chasing quick fixes or the latest trends. When you evaluate with that mindset, you're not just choosing a technology platform. You're establishing the principles that will shape how patients access care for years to come



Jeff Gartland

Chief Executive Officer
Reliant

RELATIENT

See Dash in Action

relatient.com