How Ready Is Your Clinic for AI Marketing?

The Quiet Data Gap Between AI Ambition and AI That Works, and Why It Matters This Year

Martin Mikek, MD | Co-CEO & Co-Founder, Carely Digital

Abstract

AI marketing has become the urgent conversation in every clinic boardroom. Vendors are pitching tools, boards are asking for a plan, and competitors appear to be moving. Underneath the noise, that question is going unanswered.

The evidence says the gap is wide, and in healthcare it runs on two levels. Most clinics are not yet capturing granular first-party patient data on individual interactions across the journey, the fine-grained signals AI personalization learns from. What is captured then sits in three silos: the website, the electronic health record, and the marketing database (a mailing platform, a CRM, or a marketing hub such as HubSpot). Marketing hubs can bridge the website and the database through lead capture, but the EHR sits outside that connection, and no traditional tool brings all three together. Only 6% of marketers have fully implemented AI in practice despite near-universal intent (Supermetrics, 2026), and the distance between current practice and what AI marketing requires is wider than the sector typically acknowledges.

This paper argues that AI marketing in healthcare is a maturity journey, not a purchase, and that its first stage, the Foundation stage, is where almost every clinic is still standing. Using the AI Marketing Canvas by Venkatesan and Lecinski as scaffolding, localized to healthcare, it reframes the question of which AI tool to buy into a different kind of first move.

Adoption Is Not Readiness

By early 2026, AI marketing sat near the top of almost every clinic leader's agenda. Boards expect a plan. Vendors are pitching tools. Competitors are signaling movement, not all of it real.

The numbers tell a quieter story. Only 6% of marketers have fully implemented AI despite near-universal pressure to adopt it (Supermetrics, 2026). 96% of IT leaders report AI integration, yet nearly 80% admit limited data access is holding them back (Cloudera, 2026). 61% of organizations say their data assets are not AI-ready (BCG, 2025). In Europe, 94% of healthcare providers are already using AI or planning to (IntuitionLabs, 2025). European clinics sit squarely inside the gap: high intent, low AI readiness.

Two gaps sit beneath these numbers. First, most clinics are not yet capturing first-party data on individual patient interactions at any meaningful depth, the fine-grained signals AI personalization learns from. Second, the records that do exist sit in systems never designed to talk to each other. AI built on either gap struggles. AI built on both fails.

For a clinic CEO or marketing lead, the useful first move is not to choose an AI tool. It is to locate the clinic honestly on a maturity curve, and ask what has to be true for the next stage to become possible.

AI Marketing Is a Maturity Journey, Not a Purchase

AI marketing is a staged capability, not a category of tool. The AI Marketing Canvas by Rajkumar Venkatesan of UVA Darden and Jim Lecinski of Northwestern Kellogg, now in its second edition from Stanford University Press, makes this argument explicit. It describes five progressive stages: Foundation (building a unified, governed data layer), Experimentation (running controlled AI pilots), Expansion (scaling proven use cases), Transformation (redesigning the operating model around AI), and Monetization (turning AI capability into differentiated revenue).

Each stage presupposes the one before it. An organization cannot skip from Foundation to Transformation by purchasing a more powerful tool. It cannot reach Monetization without passing through Experimentation and Expansion first. The sequence matters more than the technology at any single stage.

That framing is worth taking seriously because the alternative framing, AI as a product decision, leads clinic leadership into a familiar trap. Pick a vendor, run a pilot, report the result, move on. McKinsey's 2025 research shows what happens at scale when this pattern plays out. 78% of organizations now use AI in at least one business function, up from 55% a year earlier. Only 1% consider their AI strategies mature. The average organization runs 4.3 pilots, and only 21% reach production scale (McKinsey, 2025). BCG reached a complementary conclusion independently. The average level of marketing maturity fell by 8% between 2021 and 2024 (BCG, 2024). Adoption rose. Capability did not.

A maturity journey is harder to run than a product decision. It is also the only kind of investment that compounds. The framework locates any clinic on the curve. The next section explains why almost all of them are still at stage one.

Stage One Is Where the Bottleneck Lives

The Foundation stage is the least glamorous in the canvas and the most decisive. It is about building a clean, unified, governed data layer that AI can learn from. No dashboards, no generative content, no predictive scoring yet, just the discipline of turning every customer interaction, across every channel, into continuous, trustworthy, reusable data.

This is where the current evidence becomes unambiguous. In August 2025 Fortune reported findings from MIT's Project NANDA: 95% of generative AI pilots inside companies are failing to deliver measurable impact on profit and loss. The organizations that do achieve returns share a specific pattern. They are twice as likely to have redesigned their end-to-end data workflows before selecting a modeling technique (MIT Project NANDA, 2025, as reported in Fortune).

Informatica's CDO Insights 2025 survey tracks in the same direction. It names data quality and readiness as the single top obstacle to AI success, cited by 43% of data leaders, ahead of technical maturity and skill shortages.

None of this is surprising if the Foundation stage is taken seriously. AI is only as intelligent as the data beneath it. A model that learns from thin first-party data produces thin output. A model that learns from fragmented records produces fragmented output. Clinics running pilots on both at once discover this the hard way, usually after a quarter or two of disappointing results that nobody can quite attribute to a single root cause.

Foundation sounds generic. In healthcare, it is the least generic problem in marketing.

Why Healthcare Data Is Harder to Unify Than Anyone Expects

In healthcare, the patient journey is split across three separate silos. The website records the digital interactions before a person becomes a patient. The electronic health record captures the treatment itself. The marketing database, typically a mailing platform, a traditional CRM, or a marketing hub such as HubSpot, carries everything after treatment, from satisfaction surveys to loyalty and preventive reminders.

No traditional tool connects all three. Modern marketing hubs can stitch the website and the marketing database together through lead capture, but the EHR sits outside that connection. The EHR is the structural break at the center of the problem.

Two gaps compound on top of each other. Most clinics are not yet capturing first-party data on individual patients in enough depth to support personalization. And what they do capture does not travel across the three silos. A 2025 systematic review in the journal Information put the consequence in one line: the patient is the only entity experiencing the journey as continuous. Everyone else sees fragments (MDPI Information, 2025).

Europe has made real progress on the clinical side. The Digital Decade 2025 eHealth Indicator Study puts the EU-27 composite eHealth score at 83% in 2024, with Slovenia at 87% (European Commission, 2025). National EHR infrastructure is largely in place. What is missing, at the clinic level, is the layer above the EHR that connects pre-treatment digital activity and post-treatment engagement into one patient view.

Regulation is pushing in the same direction. The EU AI Act reaches key compliance milestones on 2 August 2026, introducing data provenance obligations that make first-party, consented data the only workable basis for AI-driven communication. The European Health Data Space, in force since 2025, gives patients direct control over their health data. Average healthcare GDPR penalties have risen to €203,000 (Secure Privacy, 2026). Clinics that do not build a continuous, consented first-party data foundation will face pressure on two fronts at once, commercially and legally.

The Hidden Prerequisite: One Continuous Patient Data Stream

The solution to the three-phase problem is not another integration project. It is a decision. A clinic that treats the patient journey as one continuous data stream has solved the Foundation stage. The stream has to be first-party, consented, and running from the first digital touch through the EHR to post-treatment engagement. It has to close both gaps at once: richer first-party capture at every step, and a single record that travels across the three silos. A clinic that keeps treating the three phases as three separate disciplines has not, no matter how many dashboards sit on top of them.

This is the prerequisite AI marketing needs. It is the one clinical leadership most often underestimates, because the conversation in the market keeps moving the other way, toward models, use cases, and generative content strategies rather than toward the data beneath them. The data foundation rarely gets airtime, which is exactly why the 95% pilot failure figure is so consistent.

A common reflex at this point is to ask whether AI itself could compensate for a weak data foundation. Could synthetic data, or foundation models trained on broader public data, fill in the gaps? The emerging evidence says no. A 2025 analysis in The Lancet Digital Health warns that synthetic data built from biased or incomplete source data reproduces the same distortions at scale. The Ada Lovelace Institute reached a similar conclusion. Scaling data volume, real or synthetic, does not rescue a structurally incomplete foundation.

The parallel precedent from consumer industries is useful here, because it removes any doubt about the order of operations. Netflix did not become the archetypal personalization business by buying AI. It spent years consolidating viewer behavior into a unified data layer, and only then let predictive and generative models run on top of it. Today roughly 80% of what is streamed on Netflix comes from AI-driven recommendations, and the recommendation system is estimated to save the company close to €0.9 billion each year through retention. Spotify's Discover Weekly, generating over 2.3 billion streams every month, was built on the same sequence. Starbucks reports a 30% lift in marketing ROI and a 15% rise in customer engagement from AI-powered predictive ordering layered on top of a unified loyalty database. Every one of these stories is a sequence, not a moment. Foundation first, AI second.

Healthcare is the industry that needs this sequence most, and the industry where following it is hardest.

What This Means for a Patient Engagement Platform

The class of tool built to operate across all three phases is the patient engagement platform. The category matters. A generic customer data platform can unify customer records in retail, banking, and travel, and does it well. Healthcare requires something more specific, because healthcare has the three-phase problem described above.

A patient engagement platform sits across all three phases by design. It provides the lead capture, content delivery, and journey orchestration on the website. It integrates into the EHR so clinical milestones can trigger the right communication without exposing clinical data to the broader marketing stack. It runs post-treatment engagement, from satisfaction requests through referral programs and preventive reminders, inside the same data environment. One patient, one record, across all three phases.

This is what makes the patient engagement platform the true first strategic purchase in a clinic's AI marketing journey. It is not the AI itself, but the substrate on which AI will eventually run. Tealium and Treasure Data both document what happens when healthcare organizations move to unified patient data platforms: better care plan adherence, lower readmission rates, stronger journey campaign performance, and more effective preventive communication. The patient engagement platform is the healthcare-specific expression of the broader customer engagement platform (CEP) category that has matured over the past decade in retail, travel, and financial services. Other industries have already understood the sequence. Healthcare has not.

After the Foundation, the Rest Gets Faster

Once the Foundation stage is built, the remaining four stages of the AI Marketing Canvas stop being abstractions and start being operational choices. Experimentation runs on full patient context rather than fragments. Expansion scales faster because the journey orchestration is already in place. Transformation gains credibility because personalized content can draw on real behavioral and clinical signals rather than demographic guesses. Monetization becomes visible because attribution runs from the first click to the completed procedure and beyond.

This is the compounding effect documented in BCG's marketing maturity research. Earlier papers in this Insights series have described the same effect as a flywheel, viewed from different angles: personalization built as a journey-wide capability rather than a campaign, marketing and patient experience unified on a single data layer so that marketing becomes the beginning of care, and attribution that moves beyond vanity metrics into first-touch, multi-touch, and eventually data-driven models.

The executive appetite for this payoff is already clear. Deloitte's 2025 Global Health Care Outlook shows that 72% of health system executives list improving consumer experience, engagement, and trust among their 2025 priorities (Deloitte, 2025a). Deloitte's Digital Transformation in Healthcare research adds that 92% of respondents name better patient experience as the top desired outcome (Deloitte, 2025b). Hyper-personalization is not a replacement for earlier stages, but the compounding effect of those stages working in sequence on a clean foundation.

The clinics that will have AI marketing five years from now are, almost without exception, the clinics that are quietly investing in their data foundation this year.

A Quieter First Move

Most clinics are being asked, right now, to do AI marketing before they are ready. The request is understandable. The market is moving, competitors appear to be moving, and a plan is expected. The problem is that the first move everyone is being pitched, the AI tool, is not the first move. It is the third or the fourth.

The first move is quieter than AI. It is the decision to unify the patient journey into a single, continuous, first-party data stream, spanning the digital funnel, the clinical record, and post-treatment engagement. That decision is less visible than a generative content platform and less exciting than a predictive model, and it is the one that compounds for the next five years.

AI marketing starts before the AI. Every subsequent stage of the AI Marketing Canvas depends on it. Every hyper-personalized experience a patient will one day have inside a clinic depends on it. Every strategic choice about differentiation, retention, referral, and preventive care that can honestly be made from marketing data depends on it.

Clinic leadership that moves first on the foundation will, paradoxically, be the leadership that moves last on AI marketing. They will be the ones for whom AI marketing, when they choose to deploy it, works.
A patient engagement platform is the shortest route to that foundation. Carely was built to be one, specifically for the three-phase healthcare problem. When the foundation is ready, the rest of the canvas becomes a set of operational choices rather than a set of open questions.

Key takeaways

- AI marketing is a staged capability, not a purchase. The AI Marketing Canvas describes five stages, and clinics are almost all still at stage one.
- Stage one is the Foundation stage: a unified, first-party, consented patient data stream. Two gaps keep clinics stuck there, thin first-party data capture and fragmentation across disconnected systems. Without the foundation, AI pilots fail.
- In healthcare, the foundation problem is structural. Patient journey data lives in three separate silos: the website, the EHR, and the marketing database.
- A patient engagement platform is the operational layer that resolves the three-phase problem. It, not the AI itself, is the first strategic purchase in a clinic's AI marketing journey.
- Once the foundation is in place, the remaining four stages of the canvas compound, each building on the one below. The clinics investing in the foundation now are the clinics that will have AI marketing five years from now.

FAQ

What is the AI Marketing Canvas, and why does it apply to healthcare?
The AI Marketing Canvas is a five-stage framework developed by Rajkumar Venkatesan of UVA Darden and Jim Lecinski of Northwestern Kellogg, now in its second edition from Stanford University Press. The five stages, Foundation, Experimentation, Expansion, Transformation, and Monetization, describe how an organization matures into an AI-driven marketing operation. It applies to healthcare because the sequence matters in every industry, and healthcare has the hardest version of the Foundation problem.

Why are most AI marketing pilots in healthcare failing?
The most common reason is not the model or the use case. It is the data foundation beneath them. MIT's Project NANDA found that 95% of generative AI pilots inside companies are failing to deliver measurable impact, and the organizations achieving returns are twice as likely to have redesigned their end-to-end data workflows before selecting the technology. In healthcare, two structural gaps make this harder than in other industries: thin first-party data capture at the individual patient level, and fragmentation across three silos (website, EHR, marketing database). AI trained on either gap underperforms. AI trained on both fails.

What does the Foundation stage look like in practice for a clinic?
In practice it means treating the patient journey as one continuous stream of first-party, consented data, starting with the first digital interaction on the website, continuing through the EHR during treatment, and extending into post-treatment engagement. It means a single patient record that three different systems and three different teams all share, rather than three different records for the same person.

How is a patient engagement platform different from a customer data platform?
A generic customer data platform unifies marketing and sales records across channels. A patient engagement platform does that, and also integrates with the EHR so the clinical journey during treatment becomes part of the same continuous record, and runs post-treatment engagement inside the same environment. It is designed specifically for the three-phase structure of the healthcare journey.

Do we need to replace our existing CRM or EHR to build this foundation?
Usually no. The EHR should remain the clinical system of record. The existing marketing tools can continue to run campaigns. The patient engagement platform sits above them, connecting the data produced at each phase into one continuous view, and can replace overlapping systems selectively rather than all at once.

How long does it realistically take to move past the Foundation stage?
It depends on the size of the clinic, the number of legacy systems involved, and the discipline of the team. A mid-sized private clinic can typically establish a working foundation and begin stage two experimentation within one to two quarters. Full maturity across all five stages is a multi-year journey, but the foundation itself is not.

Is this relevant for smaller clinics, or only for large groups?
It is especially relevant for smaller clinics. The AI that will work well for smaller clinics is applied AI rather than models trained from scratch, and applied AI runs reliably on relatively modest data volumes as long as the data is clean and continuous. The advantage of smaller scale is that the Foundation stage can be built faster, with fewer legacy systems in the way.

References

Ada Lovelace Institute. *Synthetic data, real harm*. https://www.adalovelaceinstitute.org/blog/synthetic-data-real-harm/

Boston Consulting Group (2024). *Accelerating AI-Driven Marketing Maturity*. https://www.bcg.com/publications/2024/accelerating-ai-driven-marketing-maturity

Boston Consulting Group (2025). *The Widening AI Value Gap*. https://media-publications.bcg.com/The-Widening-AI-Value-Gap-October-2025.pdf

Cloudera (2026). *AI Readiness Report*. https://www.dqchannels.com/news/cloudera-ai-readiness-report-why-96-adoption-doesnt-equal-ai-success-in-2026-11731340

Deloitte (2025a). *2025 Global Health Care Outlook*. https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2025-global-health-care-executive-outlook.html

Deloitte (2025b). *Digital Transformation in Healthcare*. https://www.deloitte.com/us/en/insights/industry/health-care/digital-transformation-in-healthcare.html

European Commission (2025). *Digital Decade 2025 eHealth Indicator Study.* https://digital-strategy.ec.europa.eu/en/library/digital-decade-2025-ehealth-indicator-study

European Commission. *EU AI Act.* https://artificialintelligenceact.eu/

European Commission. *European Health Data Space*. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space_en

Fortune (2025). *MIT report: 95% of generative AI pilots at companies failing* (reporting on MIT Project NANDA research).
https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

Informatica (2025). *CDO Insights 2025 Survey*. https://www.informatica.com/lp/cdo-insights-2025_5039.html

IntuitionLabs (2025). *AI in Private Medical Practice: 2025 Adoption Trends*. https://intuitionlabs.ai/articles/ai-adoption-private-medical-practice

The Lancet Digital Health (2025). *On synthetic data limitations in healthcare.* https://pmc.ncbi.nlm.nih.gov/articles/PMC12778113/