




**There is no “right”
AI strategy in
healthcare marketing
— *ONLY YOURS***

A guide to deciding what AI work matters for your organization.

Table of contents

Title	Page
Prologue	5
Introduction	8
Shift 1: From copying playbooks to defining your own	10
Shift 2: From “what should we do?” to “who are we?”	12
Shift 3: From aspiration to honest fluency	15
Shift 4: From shared tools to your innate advantage	17
Shift 5: From scattered effort to your next move	20
What to do with this	22



Prologue

Before you decide what to do about AI, you need to understand something most organizations miss. AI is exposing patterns that were already there.

In the pages ahead, you'll see references to two ideas: AI archetypes and AI fluency levels. They help define those patterns and explain why the same AI strategy succeeds in one organization and stalls in another.

Every organization has a default posture toward uncertainty. A pattern in how it reacts when something new, ambiguous, and high-stakes shows up. Some organizations move fast and learn by doing. Others slow down to protect trust. Some focus on building durable systems. Others prioritize proof, adoption, or long-term advantage.

You'll see those patterns surface in different ways: an Athlete pushing forward, a Skeptic slowing things down, an Optimizer looking for evidence, a Builder trying to get the foundation right. We call those patterns archetypes.

Unlike fluency, archetypes are relatively stable. They reflect how the organization instinctively responds to uncertainty, pressure, and change. The goal is not to become a different archetype. The goal is to understand the one you already are well enough to recognize which AI moves are likely to compound your strengths, and which ones are likely to create friction, resistance, or failure.

At the same time, organizations operate at very different levels of AI readiness. Not in theory, but in practice. How consistently teams can use

these tools. How well they can evaluate them. How effectively they can integrate them into real work without creating chaos, risk, or rework.

We call that fluency, and it's not static. Organizations move through levels over time. Like learning a language, early stages are often rigid and uncertain. Teams rely heavily on rules, prompts, and guardrails because they haven't yet developed the instinct to recognize what works and what doesn't. A team operating at this stage would be a Level 1.

As fluency grows, the relationship with the technology changes. Teams become more confident experimenting, evaluating nuance, recognizing mistakes, and integrating the tools naturally into everyday work. Organizations operating this way are functioning at a much higher level of fluency.

That progression matters because certain AI initiatives only work when the organization has built the fluency required to support them. Work that succeeds inside a highly fluent organization can easily stall or create chaos when copied into one that hasn't developed the same capabilities yet.

That's why lower-fluency organizations need different priorities, governance, and expectations than organizations operating at higher levels. The goal is not to imitate someone else's AI maturity. The goal is to move deliberately from your current level to the next one.

Those two forces, your archetype and your fluency, shape what AI work will actually succeed inside your organization. They determine which moves create progress, and which quietly stall.

Prologue

Most organizations never stop to define either one honestly. Vendor-driven AI strategy often skips this step because it's designed to be repeatable. The recommendations have to work everywhere, which usually means they aren't deeply grounded in the realities of any one organization.

So teams copy visible tactics from organizations operating under completely different conditions. A flashy AI initiative. A conference presentation. A vendor promising transformation in ninety days. The playbook gets copied before anyone asks whether the organization adopting it is actually built to make it work.

You already know that's not how healthcare marketing works.

As you read, you'll start to recognize these patterns in your own organization. You may notice where your instincts help you. You may notice where they limit you. You may start to see that some of the friction around AI inside healthcare organizations has less to do with the technology itself, and more to do with the mismatch between ambition, readiness, and organizational identity.

That's the point.

The goal here is simpler than that. It's to give you enough language to recognize the patterns and start asking a better question than "What should we do about AI?"

Instead ask, "What should we do next, given who we are and where we actually stand?"

Introduction

The most common question I get from healthcare marketing leaders right now is understandable but also not very useful: “What should we be doing about AI?”

The honest answer is, it depends. That can sound like a dodge. It isn't. What you should do about AI depends on three things that are mostly invisible from outside your organization and getting them wrong is the most common reason AI initiatives in healthcare marketing stall.

Who you are.

Where you sit right now.

The one advantage you have that no one else can copy.

Most AI advice in the market skips all three. It treats AI strategy as a list of tactics that work the same way everywhere. They don't. The same recommendation creates progress in one organization and waste in another, depending on the conditions it lands in.

This guide is about defining those conditions clearly enough to make better decisions. Not just about what AI could do, but about what your organization is actually positioned to do well right now, and what comes next.

Introduction

The five shifts that follow are designed to help you identify those conditions honestly. Each one names a pattern I see repeatedly inside healthcare marketing organizations. Each one points toward the adjustment that helps an organization move forward, based on its own orientation, readiness, and strengths.

The goal is simple: stop running someone else's AI strategy and start building yours.

Shift 1



From copying playbooks to defining your own

A marketing leader goes to a conference, hears a case study from a health system doing visible work in AI (Cleveland Clinic, Mayo, Northwell, take your pick), and comes back with a list.

AI-generated service line content. A custom model trained on the brand voice. An app for tracking ad inventory. LLMs in the patient experience layer.

The list becomes a roadmap. The roadmap becomes a request to the agency. Six months later it stalls. Or it works in the demo and dies in the org.

Shift 1

When I dig into one of these, the diagnosis is rarely that the tactics were wrong. The tactics were probably fine on their own. The diagnosis is that the playbook belonged to a different organization. The system that built it had a different posture toward risk, a different relationship between marketing and compliance, a different bench of internal expertise, a different definition of what success looks like a year out. You imported the moves without importing any of the conditions that made them work.

The pattern is familiar from other parts of the work. We don't tell every hospital to run the Mayo Clinic patient experience playbook, because Mayo's playbook is built on Mayo's identity, governance, capital, and culture. We don't tell every system to copy HCA's brand approach, for the same reason.

Why we let ourselves do it with AI is partly the novelty, partly the pressure to show progress, and partly the way vendors talk. The move is the same one we'd correct in any other strategy meeting.

The shift / Stop benchmarking your AI strategy against systems whose orientation, fluency, and strategic position you don't share. Benchmark against an honest read of yourself. The next four shifts are how you do that.

Shift 2



From “what
should we do”
to “who are we?”

**Every organization reacts to change in its own way.
AI exposes that pattern fast.**

Some teams move to learn. Some move to protect. Some move only when the evidence is in. These tendencies already exist in your organization, whether anyone has named them out loud or not. The work is to name yours and choose AI moves that match.

Here’s an example.

Shift 2

Imagine two hospital marketing departments that look similar on paper. Same bed count, same competitive set, same service line emphasis.

One has tendencies we'd call an Athlete orientation: wired to compete, perform, test in public, learn fast, treat market position as something you train for. The other has a Skeptic orientation: wired to scrutinize claims, demand evidence, protect the institution from the cost of moving on something that turns out to be hype.

Hand both teams the same recommendation, "Use generative AI to scale personalized service line content."

For the Athlete, that's a green light. But only if the team is operating at a fluency that lets them tell good output from bad fast enough to keep up with their own ambition. Without that, the same green light produces a content factory that quietly degrades the brand for two quarters before anyone notices.

For the Skeptic, the same recommendation is a non-starter until the system can answer how do we know this is safe, accurate, on-brand, defensible. That answer doesn't come from a vendor deck. The Skeptic's correct first move with AI is rarely tactical. It's building the evidence apparatus that lets the rest of the organization trust the move when it comes.

Same recommendation. Two right answers. Neither team needs to become the other one. Each team needs to lead with its own posture and address its own blind spot. The same logic produces different right answers for

Shift 2

an Optimizer (whose superpower is measurement and iteration) versus a Visionary (whose superpower is seeing where the puck is going). The recommendation that fits one will pull the other off its edge.

The shift

Before scoping any AI initiative, force the conversation about who you are. What kind of organization, what kind of marketing function inside it, and how does that shape what “right” looks like for you?

Shift 3



From aspiration to honest fluency

The second question is where you are right now. Your real, on-the-ground familiarity with how these tools behave, how to evaluate them, how to govern them, how to fold them into the work.

Fluency is a capability your organization either has or doesn't have. The moves that pay off at low fluency are different from the moves that pay off at high fluency, even for the same orientation.

Shift 3

The trap is self-assessment that isn't honest. The Athlete who thinks her team is at fluency level 4 when it's at 2. The Skeptic who's calling level 1 caution "rigor." The Optimizer who has been told he needs to be more visionary and is now drifting away from his edge.

When fluency gets inflated on paper, the next AI investment lands wrong. You scope work for a team that doesn't yet exist, then call it execution failure when it doesn't deliver.

There's a structural reason this is hard. What your team is best at today requires fluency it hasn't yet built to become a strength in the AI context. Generating that fluency requires work the team hasn't started. The work feels uncomfortable to start because the team is currently operating from its existing strengths. So most leaders move sideways instead of forward, picking up the visible artifacts of someone else's progress without doing the unglamorous work that earned them.

The shift / Get a real read on where you are, not where you wish you were. Most of the AI failure I see in healthcare marketing is downstream of self-assessment that wasn't honest.

Shift 4



From shared tools to your innate advantage

This is the shift teams skip the most. It's also the one with the highest ceiling.

There are two kinds of AI use, and most strategy conversations conflate them.

The first kind is universal. AI helps your team draft faster, summarize meetings, run a QA pass, generate a creative concept. The time saved is real. The value is real. But the tools doing that work are available to every system, every agency, every competitor in your market. Whatever ground you gain, others gain too. The advantage caps at how well you keep up.

Shift 4

The second kind is specific. AI applied to something only your organization can do. Your data. Your patients. Your physicians. Your campaign history. Your community presence. The advantage here doesn't cap, because no competitor has access to the inputs.

Most teams spend almost all their AI energy on the first kind, treat the second as something to get to later, and then wonder why their AI strategy looks like everyone else's.

The shift isn't away from universal use cases. Keep those. They pay for themselves. The shift is toward an honest inventory of what your organization has that no competitor does, and what AI lets you finally do with it.

Picture a regional system anchored in South Louisiana. Eight hospitals, dominant in Acadiana, decades of relationships with families seeing the same physicians for two generations.

That system has assets a national vendor or an out-of-market entrant cannot replicate. Twenty years of service line campaign data, with results broken down by parish, by zip code, by month, by creative approach. Hundreds of thousands of patient communications (call center transcripts, MyChart messages, post-visit surveys) in the language patients in the region use about their concerns. Physician panels with multi-generational trust and outcome data that compounds over time. Specific knowledge of which employers in the region are on which carriers, and which contracts are about to turn over.

Each of those is dormant strategic value. Each was previously too expensive or too slow to put to use at scale. AI changes the cost structure.

Shift 4

Of course, not all of these are usable assets given the sensitivity of the data therein. But that's beside the point. Start paying attention to what you have that can't be replicated easily.

Real Type 2 work for that system might look like a tool that lets a service line marketing manager ask, "What's worked for breast cancer screening in 50 to 65 year old women in our southern parishes over the last decade?" and get an answer grounded in evidence. It might look like patient communications shaped by the way patients in the community talk, not generic hospital voice. It might look like proactive outreach to HR teams at regional employers whose carrier contracts are shifting, before competitors notice.

A hospital system in Boston cannot do this for that market. A national agency cannot do this. A generic AI tool cannot do this. Only the system itself can. That's where strategic AI investment compounds.

The shift / Before scoping any AI initiative, ask whether it would still be a differentiator if your biggest competitor stood up tomorrow with the same AI tools you have. If not, you're working in the universal lane. Treat it as table stakes, not strategy. If yes, because the work draws on something only your organization has, that's where strategic AI investment belongs.

Shift 5



From scattered effort to your next move

Once you know who you are, where you are, and what you have, the next move usually announces itself.

What gets in the way is everything else. The pace of new tool releases. The competitor announcement. The board member who saw something at a conference. The vendor pitch that lands in your inbox the same week.

Each input is reasonable on its own. Together they pull teams in directions that don't compound. Marketing tries content velocity while the contact center pilots summarization while a service line tests personalization. Three

Shift 5

efforts, three vendors, three definitions of success. None of them tied to a decision about who you are or where your real advantage sits.

The shift is to narrow the field. Decide what continues, what stops, and what scales. Capture what works. Retire what doesn't. Reduce the number of active efforts so progress can build on a defined set of decisions instead of getting diluted across all of them.

What this looks like depends on the answers you got in shifts two through four. An Athlete at fluency 3 with a clear proprietary advantage around physician engagement should be running a small number of well-instrumented experiments in that lane and almost nothing else. A Skeptic at fluency 1 with a clear advantage in patient communications should be building the governance and evaluation muscle that makes the eventual move defensible, not running parallel pilots for the sake of looking busy.

The next move that's right for you is the one that fits what you've defined about yourself, which may not be the move being flaunted in a keynote.

The shift / Run the AI plays that match your orientation and fluency, then lean on the one advantage no one else has. The rest is noise.

What to do with this

The work is to define who you are, where you stand, and the work only you can do.

Once those three answers are honest and named, your AI strategy stops being a question. The next ninety days get specific. The choices you don't have to make become as important as the ones you do.

If you'd like a structured way to answer the first two questions, we built a model for it. The AI Archetype and Fluency Model is a short assessment that produces a roadmap unique to your team. Find out you're an Integrator at fluency level 3, and the AI strategy gets clear. Your only AI goal is moving from Integrator L3 to Integrator L4.

If you'd like to take the assessment, reach out.

**UN
LOCK** 