



WHITE PAPER · AI TRANSFORMATION

People First, Then AI

Why the teams doing the work are what make AI move faster

Prepared for: Healthcare and financial services executives leading AI programs
Prepared by: Jim Barto, SVP Solutions, Curate Partners | October 2026

Involving the people who do the work is what speeds AI up

Most large organizations have more AI ideas than they can act on. What they lack is a reliable way to turn those ideas into capabilities their teams trust, use, and depend on. The market offers two answers. The first is a large-scale rebuild that asks for a multi-year commitment before value appears. The second is a long list of disconnected pilots that rarely reach production and are hard to tie to a return.

We take a different path. We work one department at a time, where the workflow is specific enough to improve quickly, and we start with the people who do the work. That choice is often seen as the slow option. In our experience it is the fast one. In our experience, programs that skip discovery do not move faster. Gaps often surface later in testing, adoption, or production, where they are more difficult to address.

This paper explains where that view comes from, what we are seeing in our AI delivery work, and how our operating model, ATOM, puts it into practice. Our central argument is simple: involving the people who do the work is what makes AI move faster.

Enfranchisement is the key to acceleration.

THE PROBLEM

Two ways AI programs stall

Across the healthcare and financial services organizations we work with, AI programs tend to stall in one of two ways.

The big bang

Some firms are sold a full rebuild: rewire the ecosystem, modernize every platform, then apply AI on top against a central architecture. The logic is sound on paper, but value arrives late, often a year or more after funding. Leadership patience runs out before the return shows up, and the people expected to adopt the result were rarely consulted while it was being designed.

Death by a million cuts

Other firms go the opposite direction. They collect individual ideas from across the enterprise and attack them one at a time. Each pilot looks reasonable in isolation, but together they are hard to manage, hard to fund consistently, and nearly impossible to tie to a return. Many live their whole lives as science projects.

The middle path: one department at a time

We work at the department level: a claims and service operations team, a client onboarding group, an advisor services function, a finance or HR department. At that level the workflow has enough specificity that improvements can be designed, tested, and measured quickly. The people doing the work know it in detail. And the results are visible enough to build confidence for the next department.

From there, a structured, templated set of steps moves each idea from concept to prototype, to a working proof of concept, to a production capability the enterprise can rely on.

“AI will change everything, so today’s process doesn’t matter”

A popular view holds that there is no need to understand how work gets done today, because AI will change it anyway. We hear it in industry commentary and in project kickoffs. We disagree.

We are a practice of people who spent most of our careers inside operating companies before becoming advisors. One lesson repeats: old technology dies hard, and old processes die hard. The reason is people. Teams build habits, sub-processes, and workarounds around the systems they use. The cost of ripping all of that out at once is usually more disruption than a business can absorb.

Starting from “don’t worry about how you do it today” creates a second problem. People who feel their work is being redesigned around them, rather than with them, tend to slow it down, sometimes without meaning to. That is not the culture most leaders want to build, and it is not a path to speed.

A better starting point is to ask how the work is done today, what skills and roles support it, what governance surrounds it, and what the people doing it are trying to achieve. That is where the near-term opportunities are found: experiments that can run now, without waiting months for a platform rebuild.

OUR INFLUENCES

What the digital revolution taught us

Many of us came up through the digital and mobile era, putting new capabilities in the hands of employees, clients, members, and consumers. Those years shaped how we approach AI today.

In healthcare

We know that if a patient schedules a preventive or diagnostic appointment, there is a better chance of heading off a serious and costly health event later. Before digital, getting people to take that step meant phone calls, mail, and sometimes in-person outreach. Digital let organizations reach people where they already were and make the next step easy, alongside the outreach that was already in place.

In financial services

Onboarding a new client into a wealth management practice once meant spreadsheets, email, and faxed documents to collect compliance, asset, and account information. Digital onboarding shortened the cycle, but the bigger gain was quality. Smart checks could flag a missing or misread answer at intake, before it created problems for the client later.

The lesson that carried over

The common thread was human-centered design. When we look back at digital products and initiatives that failed, a recurring root cause was a skipped or shortened discovery phase. Someone was confident they already knew the answer. Talking to users, mapping the current journey, and agreeing on the jobs to be done felt like delay. Teams went straight to requirements, or straight to building.

Those projects did not go faster. Problems surfaced late in testing, or worse, products reached the market in search of a problem to solve, which is the fastest way to lose credibility with the people you are building for. Many of us learned this the hard way, more than once.

We see the same pattern in AI. Teams are encouraged to skip to the build stage. But an agent cannot simply be pointed at a whole business unit and told to automate it. Someone has to decide where in the process AI makes sense, and that decision depends on understanding the work and the people doing it.

WHAT WE ARE SEEING

Lessons from our AI delivery work

Our practice has delivered more than 50 AI use cases across functions including sales, marketing, legal and contracting, customer service, engineering, clinical and operational analytics, provider operations, finance, and HR. Several lessons hold across that work.

~10% of proposed use cases did not need AI at all. They were process, data, or training problems.	~30% could be solved with tools the client already owned. No custom build required.	~98% reduction in manual testing effort when a stalled prototype gained a structured validation method.
---	---	---

Discovery changes the answer

For a large integrated payer and health system, structured intake across more than a dozen business areas showed that roughly 10 percent of submitted use cases did not need AI, and about 30 percent could be addressed with enterprise tools already in place. Without discovery, both groups would have consumed build budget.

Discovery also reshaped the use cases that did proceed. An HR team expected a simple benefits assistant, but real value required personalized answers tied to each employee's elections, which meant integration and access controls. A contracting team expected AI to cut contract cycle time, but legal review steps dominated the timeline. The near-term win was fewer errors and less rework, with intake reduced from days to minutes. A marketing team wanted personalized content at scale but did not yet have the customer data to support it, so the first release was scoped to fit what was ready.

The people closest to the work hold the hidden logic

In a utilization reporting engagement, one operations leader spent eight to ten hours a week reconciling timecards, rates, and project plans in a workbook. When our consultant shadowed the routine, steps described as a short sequence turned out to contain many embedded checks and judgment calls. Automating from a verbal description alone would have dropped them. The process owner's knowledge became the design input, and her role shifted from assembling data to governing it.

Value is more capacity, not fewer people

The question we hear most is whether AI lets an organization hire fewer people. In our experience that is rarely the right goal. The better question is how much more a strong team can do with what it already

knows: higher throughput, better quality, fewer errors, faster response. The people doing the work are the subject matter experts. The aim is to get more out of that expertise.

Start with what you already have

Many organizations already own capable AI tools. Microsoft Copilot, for example, provides access to leading models and connects to enterprise content in SharePoint. A well-designed proof of concept built on existing tools can show whether an idea has value before larger investment is committed. The test is not whether a tool is impressive. It is whether it moves a measure the business cares about.

Set clear limits on what AI decides

In the utilization engagement, a Copilot agent answers factual questions about project status, but it is deliberately limited to read-only retrieval. It does not advise on staffing or client decisions. That boundary is part of why the organization trusted it. AI retrieves the facts and accountable people make the decisions.

Validation is what gets prototypes to production

A healthcare organization had an AI pricing assistant that showed promise but lacked a consistent way to test it. Adding a structured validation framework with reusable test cases cut manual testing effort by about 98 percent and moved the assistant into production, with a repeatable model for the next use case.

Reuse is what makes scale affordable

In larger programs, a core team working ahead of delivery teams can build shared components once for reuse. In one program, a single data connector served more than ten use cases. That is the difference between a coordinated program and a collection of separate projects.

OUR APPROACH

How ATOM puts people first

The Curate AI Accelerator is built on ATOM, our AI Transformation Operating Model. ATOM organizes AI work across four pillars, and the order is deliberate.

PILLAR	WHAT IT COVERS
People	Skills, roles, and the governance culture that make adoption real. This is where every engagement starts.
Process	How ideas are qualified, governed, built, and proven.
Platforms	The shared technical foundation and tooling solutions are built on.
Projects	The portfolio of use cases where value is actually created.

A hypothesis before a build

“That’s a cool AI idea” is not a reason to pursue something. Each idea is stated as a testable hypothesis: based on what we observe, we believe that if we do X, Y will improve, and we will measure it this way. Teams sometimes want to skip this step. It is the step that makes the return provable later.

Nine stages, two hard gates

ATOM moves each use case through nine stages. The early stages carry the most weight, because most of the value comes from the discipline applied before anything is built.

STAGE	FOCUS
1. Discovery	Understand the problem and the people doing the work
2. Qualification	Decide how it should be solved, including whether AI is needed
3. Baselining	Agree the value to beat and the measure that will move
4. Feasibility	Confirm it can be done through a low-cost experiment
5. Funding (hard gate)	No build without a business case and governance approval
6. Design	Specify the durable solution
7. Mobilization	Stand the work up
8. Validation (hard gate)	No launch without formal proof that it behaves
9. Operation	Run it and confirm the value continues

Governance that moves with delivery

Every delivery milestone has a matching governance checkpoint: registering and risk-scoring the request, clearing data, privacy, and vendor terms, setting oversight rules, approving funding and launch, and monitoring behavior once live. The goal is not more control. It is fewer surprises and a faster path for work that is already sound.

Run each capability as a product

Once live, each capability is measured against the value it was meant to deliver and improved over time. In effect, the process turns the people closest to the work into product managers for the tools that support them. That ownership is what keeps AI capabilities in use after the project team moves on.

Method creates speed

None of this has to be slow. The structure exists to organize thinking, decisions, and work so teams can move quickly with confidence. What most organizations lack is not the will but the background and bandwidth to run it consistently. That gap is why work drifts toward science projects at one end or multi-year transformations at the other.

WHERE IT APPLIES

Examples from our work

The approach applies wherever a department has a defined workflow and people who know it well. Some examples apply across industries.

AREA	EXAMPLE	HOW PEOPLE STAY IN THE LOOP
Sales	Call summaries logged to CRM; account and prospect research; proposal and RFP pre-population	Reps review and approve before anything is saved or sent
Marketing	Brand-aligned content drafts; compliance review of assets; visibility across AI answer engines, not just the website	Writers, brand, and legal reviewers refine and approve
Legal and contracting	Guided contract intake; obligation extraction into tracked task lists	Mandatory legal review stays in place
Customer service	Call summarization; real-time knowledge suggestions; guided flows	Agents confirm notes and handle complex cases with 'warm hug' familiarity
Analytics	Plain-language questions against operational and financial data	Analysts validate results before they are shared

Industry-specific examples, such as utilization management and service operations in healthcare, or advisor experience and client onboarding in financial services, are covered in our companion briefs for each industry.

FOR LEADERS

What this means for AI leaders

- **Pick a department, not the enterprise or a single task.** Specific workflows move faster and prove value sooner.
- **Do not skip discovery.** It is the step most often cut and the one most often behind failed initiatives.
- **Make every idea earn its funding.** State the hypothesis, agree the baseline, and test cheaply first.
- **Measure capacity, not just headcount.** Throughput, quality, and response time often matter more.

- **Build governance in, not on.** Checkpoints that move with delivery prevent late rework.
- **Treat AI capabilities as products.** Give the people closest to the work ownership of what they use.

The organizations that move fastest are the ones that bring their people with them.

ABOUT CURATE PARTNERS

Who we are

Curate Partners is a practitioner-led consulting firm working in data, analytics, AI, digital strategy, and operating model design, primarily for regulated healthcare and financial services organizations. Many of our people led digital and data work inside operating companies before becoming advisors, and that experience shapes how we work with clients.

Curate has always been a people-first organization. Our heritage is in finding, placing, and building teams around talented people, and that instinct carries into every engagement. We lead with the people and land with the framework.

We support clients across the full path from strategy to production:

- **Advisory:** AI strategy, operating model design, and governance.
- **Enterprise data:** Data architecture, master data management, semantic layers, and modern data platforms such as Databricks and Snowflake. Data in context is the foundation every AI use case depends on.
- **Delivery:** Teams that build and run the work, including more than 150 engineers currently deployed at a Fortune 10 enterprise client.
- **Talent:** A talent solutions business that helps clients build the teams they need.

The Curate AI Accelerator, powered by ATOM, brings these together into a repeatable path from AI idea to measurable value.

NEXT STEP

To discuss how this approach applies to your organization, contact Jim Bartro, SVP Solutions, at jim.bartro@curatepartners.com | curatepartners.com