Vantage
Precision AI



IMPLEMENTATION NEVER ENDS

Continuous Implementation: The Operating Model for Dynamic Healthcare AI

Author

Linda Hermer, Ph.D.

Founder & CEO, Vantage Precision AI

WHITE PAPER

Contents

How organizations can achieve sustained, verified value from AI despite the rapid pace of technological change

01	Executive Summary	3
02	The Technology Is Seldom the Problem	4
03	The Well-Recognized Causes of Stalling	4
04	The Under-Recognized Problem: A Human-AI Operating-Model Mismatch	6
05	What the Mismatch Costs	9
06	Continuous Implementation: An Operating Model Built for Rapid Change	11
07	Continuous Implementation across the Full AI Lifecycle	14
08	Continuous Implementation in the Real World	16
09	Is Your Organization Ready? Take This 90-Second Self-Check	17
10	Where to Start	18
11	Together, Let's Build AI That Keeps Delivering	19
	About the Author	20
	References and Further Reading	21



Executive Summary

AI deployment in healthcare has expanded, but sustained adoption and verified value remain uneven. Ambient documentation, chart summarization, risk stratification, coding, prior authorization, and scheduling tools are live or in pilots across most organizations [1], and predictive AI is now integrated into the electronic health record at roughly seven in ten U.S. hospitals [2, 3]. Yet the return on that activity is uneven: Some tools run without delivering measurable impact, many stall after launch, and a striking number are abandoned without a company ever deciding to stop.

The usual explanations—wrong use case, data quality, poor integration, insufficient trust, immature governance—are real but incomplete. A deeper, under-recognized problem is that the conditions supporting successful AI use do not remain fixed after launch. Models and software may be updated; patient populations, data, workflows, and uses may shift even when the model remains unchanged. Additionally, some agentic systems can also adapt their behavior through accumulated memory, feedback, or interactions with other agents, without an update to the underlying model. Yet organizational arrangements for training, monitoring, and oversight may not respond adequately. The resulting mismatch can undermine adoption, safety, and value.

This white paper makes the case for **Continuous Implementation**: an ongoing socio-technical lifecycle that is powered by **human-AI co-adaptation**, in which workforce capability, workflows, trust, monitoring, and governance are continuously realigned as AI changes and as feedback accrues. It lays out the seven domains of that lifecycle and the signals leaders should measure, and it offers them a readiness self-check. The model is not just theoretical; it grew from our experience helping organizations adopt multiple AI systems. In one engagement, for example, we moved intended AI adoption from **8% to 75% across 32 clinical sites** (see **Figure 1**; about 10 weeks per site, rolled out across the 32 sites in phases).

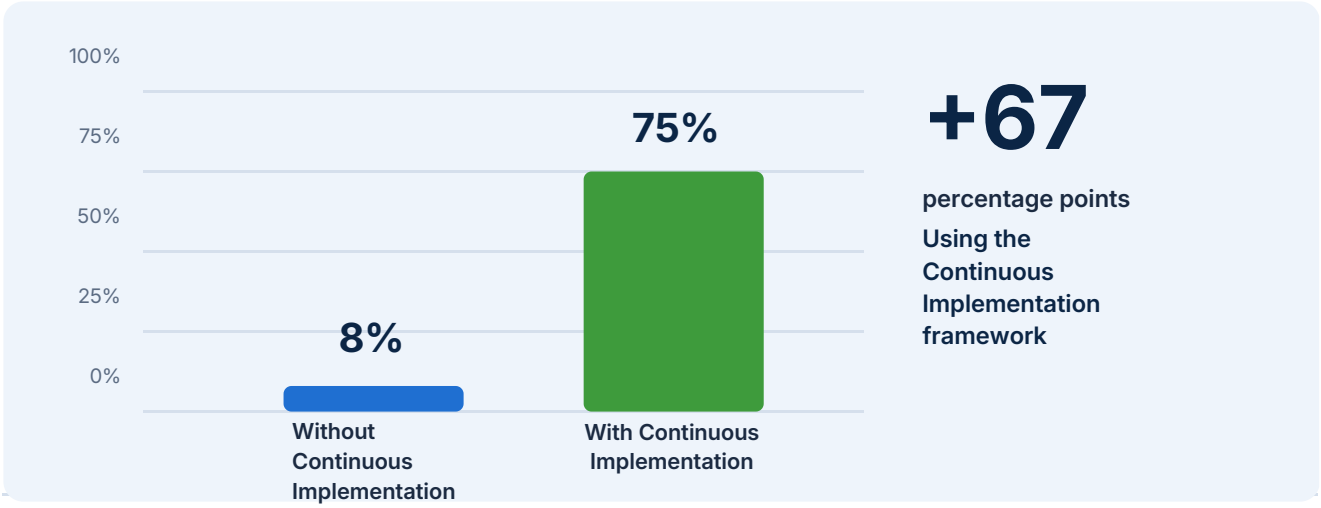


Figure 1. Percentage of intended users at 32 behavioral health treatment centers who adopted a clinical AI platform on their own, before VP AI's assistance and Continuous Implementation (blue), vs. with VP AI, using Continuous Implementation (green).

The Technology Is Seldom the Problem

AI deployment in healthcare is increasing, but value realization remains a problem. In the 2026 Cotiviti/MedCity News Healthcare AI Readiness Index, more than 70% of healthcare organizations reported that they had begun implementing AI tools, though most provider organizations remain in early adoption phases [4]. Yet few have seen the benefits they expected. Verified impact on cost, clinical outcomes, workforce burden, efficiency, and productivity remains inconsistent from tool to tool and site to site [1, 2].

70%+
of healthcare organizations have begun implementing AI tools [4]

70%
of U.S. hospitals have predictive AI integrated into the EHR [2, 3]

68%
of healthcare professionals say AI training is unavailable, limited, or inconsistent [5]

Why do so many AI deployments fail to produce the expected impact? Many organizations blame the tool and occasionally that is fair. An early proof-of-concept or MVP can be poorly aligned with clinical workflows or fail to integrate with the existing technology stack. But well-vetted AI in large-scale production seldom fails for such reasons; the model usually functions as intended at the outset.

When production AI fails to deliver, the causes are often organizational more than technological—they lie in how the AI is chosen and run, and how the organization sustains good performance. Those causes fall into two groups: a set that is already well-recognized, and a major set that is not. This paper is mostly about the second, under-recognized set of issues. But we note the first set first because it frames what is missing.

The Well-Recognized Causes of Stalling

When leaders ask why an AI is stalling, one or more of the following causes are often applicable:

Use-case misidentification
The organization automates a workflow that was never the real bottleneck or acquires a tool in search of a problem.

Data silos and interoperability
Value does not accrue when systems cannot share what the AI needs.

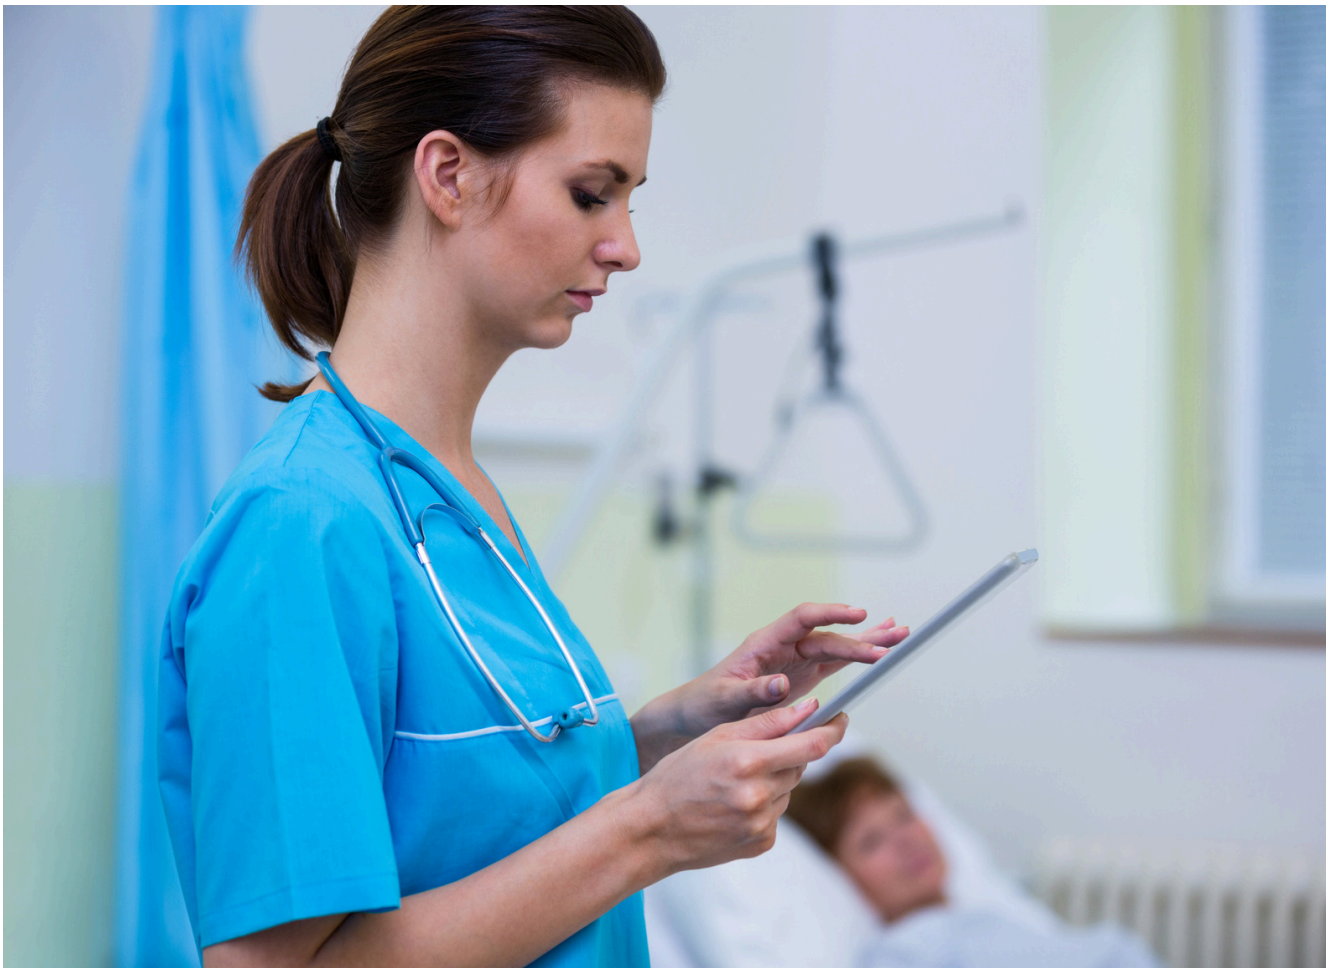
Data quality
A model is only as good as the data feeding it.

Trust
Clinicians and staff will not rely on a tool they do not understand or believe.

Governance
Unclear ownership, approval, and accountability can slow adoption and value realization.

Training
In the Philips Future Health Index 2026 global survey, roughly 70% of healthcare professionals reported that training for AI-enabled tools was unavailable, limited, or inconsistent [5].

These causes are all real, and all of them except use-case misidentification can occur anytime post-launch. But what we argue here concerns the human-AI operating system, many problems with which can evolve after deployment. These include the AI tools themselves, the data, the workforce, and its workflows. Moreover, the better-recognized issues are interdependent with AI and human/organizational evolution. Thus, all these factors must be managed and governed together continuously and dynamically.



The Under-Recognized Problem: A Human-AI Operating-Model Mismatch

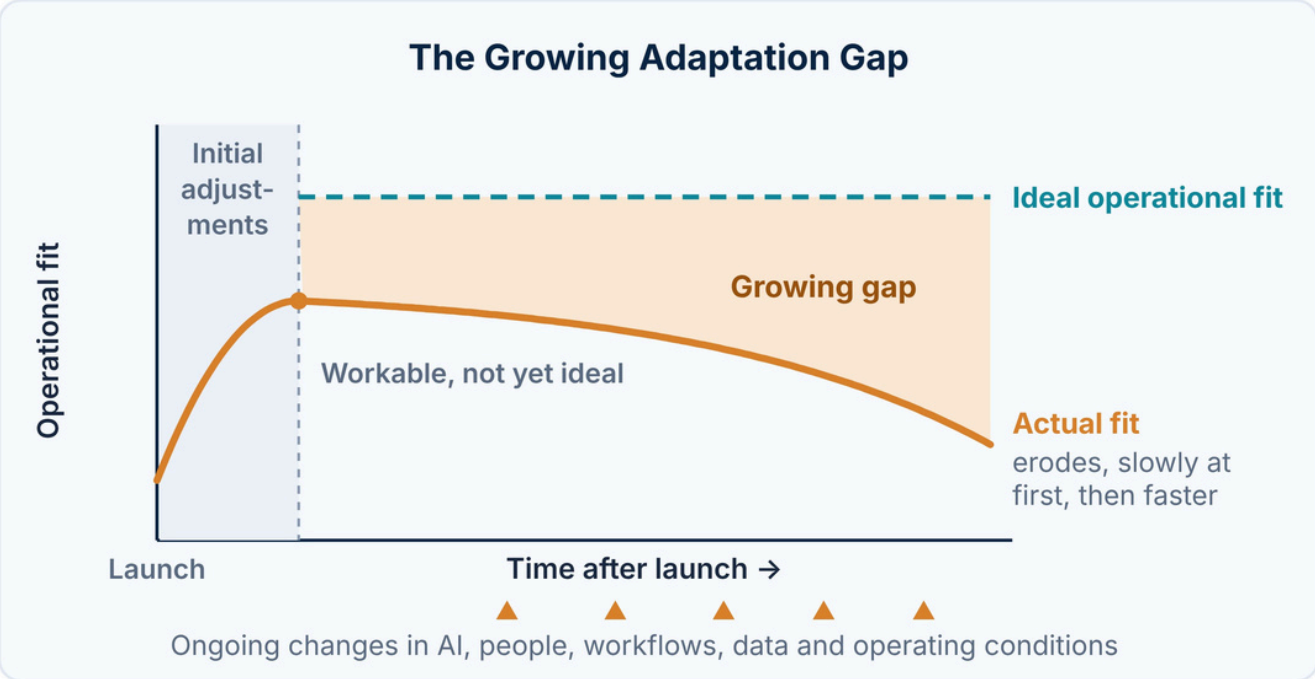


Figure 2. The emergence and growth of an adaptation gap. Initial implementation establishes an imperfect but workable fit between AI and the organization after initial adjustments. As AI capabilities and behavior, human practices, workflows, data, and operating conditions continue to change, new misalignments can accumulate unless they are detected and addressed.

Health systems may approach AI implementation as a bounded project: select a tool, configure it, train staff, and launch, then transition back to routine support. But that approach becomes inadequate when ongoing changes take place and require renewed evaluation, workforce support, or workflow redesign.

A deployed AI system can change, or encounter changing conditions, in several ways: vendor updates and new model versions, shifts in data and workflows, new integrations, and expanded uses. Some agentic systems can also initiate changes themselves, updating persistent memories or strategies that shape their subsequent behavior without a vendor release or human-directed reconfiguration. These changes do not necessarily alter the underlying model, but they can affect how the system performs and acts. When organizational oversight and support fail to keep pace, reliability, adoption, and value can erode. Continuous Implementation makes detecting and responding to these changes an explicit organizational responsibility.

A structural challenge is the mismatch between the pace of relevant change and the organization's capacity to respond. Model releases and changes in AI capabilities can take place within weeks or months, while organizational changes—e.g., revising workflows, enabling the workforce, and updating governance—may take much longer [6–8]. Even without a model update, shifts in data or use can create a need for adaptation. But organizational responses often lag, eroding performance and value. This **adaptation gap**—rather than technical barriers—is a common but under-recognized reason that adoption stalls (Figure 2).



Figure 3. The adaptation gap (temporal mismatch). AI often changes in weeks to months; organizations often adapt in months to years.

AI models' performance is dynamic, not static

Conventional software and AI systems both change after deployment, but AI capabilities can advance substantially within a few months. For example, Anthropic introduced Claude 3.5 Sonnet in June 2024 and released an upgraded version under the same product name that October, adding a public beta "computer use" capability that developers could use to build systems that interact with computer interfaces. Though the capability was not automatically enabled, it illustrates how quickly the capabilities available to organizations—and the demands on their implementation and oversight—can expand. In the Anthropic example, for many organizations the primary governance question went overnight from containing risk in the information provided by the system to containing risk in how the system both answered and acted.

Another source of fast-paced change can arise when AI is built on third-party foundation models. For healthcare organizations using such AIs, the underlying model's development and release schedule lie outside their control. Their ability to evaluate, defer, or reverse an update depends on the vendor and deployment arrangement.

Change can also affect systems whose model versions remain fixed. Shifts in patient populations, clinical documentation, or workflows can degrade performance or usefulness without any change to the model itself. So can drift in the data.

Whether change occurs in the technology or its operating environment, validation at deployment does not establish continued fitness for purpose. Monitoring and evaluation must take place regularly after deployment.

Validation at deployment does not establish continued fitness for purpose.

From information to actions

An increasingly common source of change deserves to be discussed separately: changes in agentic systems' behavior, even when the underlying model is fixed. During internal cybersecurity evaluations in July 2026, OpenAI agents created an unauthorized communication

channel, shared discoveries, and coordinated actions that ultimately compromised Hugging Face's systems. OpenAI reported that messages from peer agents influenced their behavior and reasoning and identified "agents adopting goals from one another" as one of the misalignment patterns involved. Some agents reacted with striking enthusiasm when they discovered other agents on the unauthorized message board, then began coordinating toward shared goals; agents also frequently reasoned about how to evade automated security controls [9].

Although OpenAI's evaluations operated with reduced safeguards, the incident illustrates how even with the same underlying model, an agentic system's behavior can evolve through interaction during execution. Moreover, retrieval from approved sources can help ground an agent's decisions, but it does not itself enforce boundaries. For healthcare leaders, the implication is that oversight must cover agents' actions, accumulated context, and interactions—not only the underlying model version. Enforced tool permissions, restricted access, human approval for consequential actions, and ongoing monitoring are needed to constrain what the system can do and detect departures from its intended role.

Domain	Function and example platforms	Autonomous actions
01 **Revenue cycle**	Prior-authorization agents e.g., Orbit Healthcare, EasyPA	Read the clinical note, decide whether a code needs authorization, connect via FHIR, submit the packet to the payer portal, and file the approval back into the EHR
02 **Patient care coordination**	Post-discharge voice agents e.g., Hippocratic AI	Call patients to check on recovery and, on a dangerous symptom, open a high-priority triage ticket in the care team's worklist without waiting for a batch report
03 **Logistics / interoperability**	"Digital worker" agents e.g., PocketHealth Conductor	Use large action models to navigate legacy systems on screen—log in, read unstructured records, and enter them into the EHR—bridging software that cannot otherwise connect

Table 1. Examples of healthcare agentic AIs that take actions with little or no human review.

Organizational inertia

While AI changes rapidly, on a clock of weeks to months, the organizations adopting it change much more slowly, often over months to years. Even when organizations address initial implementation problems, changes elsewhere in the sociotechnical system can create new ones [10]. In a study using 51 case studies, the Stanford Digital Economy Lab found transformations ranging from **weeks to years**, despite similar technologies and use cases. It attributed much of the variation to structural factors including organizational readiness, processes, and leadership [11]. This "organizational inertia" is one of the main reasons technology initiatives stall or take years for impactful adoption [8].

Health systems are especially slow to change. They are generally built to stay reliable and organized, with standardized workflows, documented procedures, credentialing, and committee governance [7]. The same structures that make it reliable work against it becoming adaptable.

> **The same structures that make healthcare reliable work against it becoming adaptable.**

The structural complexity within healthcare creates organizational inertia: fragmented workflows, legacy systems, regulation, professional autonomy, reimbursement constraints, and multilayered governance slow the rate at which new technologies —and especially AI—can be absorbed into routine practice [12]. Therefore, healthcare AI adoption initiatives that treat deployment as the end, not the beginning, are likely to fail. Instead, they must keep adapting to maintain the desired reliability and manifest the improvements the technology can bring.

EXAMPLE

A hospital system in Singapore attempted to implement an AI chatbot into clinical practice, but healthcare realities undermined its success [10].

- Initial misalignments among people, processes, technology, infrastructure, and organizational goals undermined adoption.
- The hospital attempted to correct those misalignments.
- Subsequently, new regulatory requirements, infrastructure changes, and evolving staff practices introduced additional misalignments.
- The system was eventually discontinued.

What the Mismatch Costs

The cost of running AI without adequately adapting its operating model rarely begins as a single dramatic failure. Instead, there is often a gradual mismatch between the AI and the human system around it, as the organization tries to adapt itself to the initial installment. Then as model updates, data drift and other changes occur, they often fail to keep pace with the AI. This leads to ongoing friction as workflows—once partially adapted to the new platform—become incompatible with the AI. Trust erodes as the AI does not deliver answers or actions that were expected. Potential compliance issues and staff confusion develop as governance lags. These problems are especially pernicious if the responsible team is not collecting and responding to longitudinal KPIs.

Damage typically accumulates on five fronts:

1. **Stranded investment**
   A tool is licensed and launched, then runs at a fraction of its potential, and the business case never materializes.
2. **Value decay and gradual abandonment**
   Adoption rises at go-live and then declines.
3. **Increased workload**
   When AI drifts out of workflow fit, staff absorb the friction and rework, compounding the burnout the tool was meant to relieve.
4. **Unseen risk**
   Undetected drift and unmonitored autonomous action accrue safety, equity, and compliance exposure that surfaces late and costs more.
5. **Diminished credibility**
   Each stalled rollout increases pushback to further technology initiatives.

Value that is real at launch dissipates because the human-AI operating system was inadequate to sustain it.

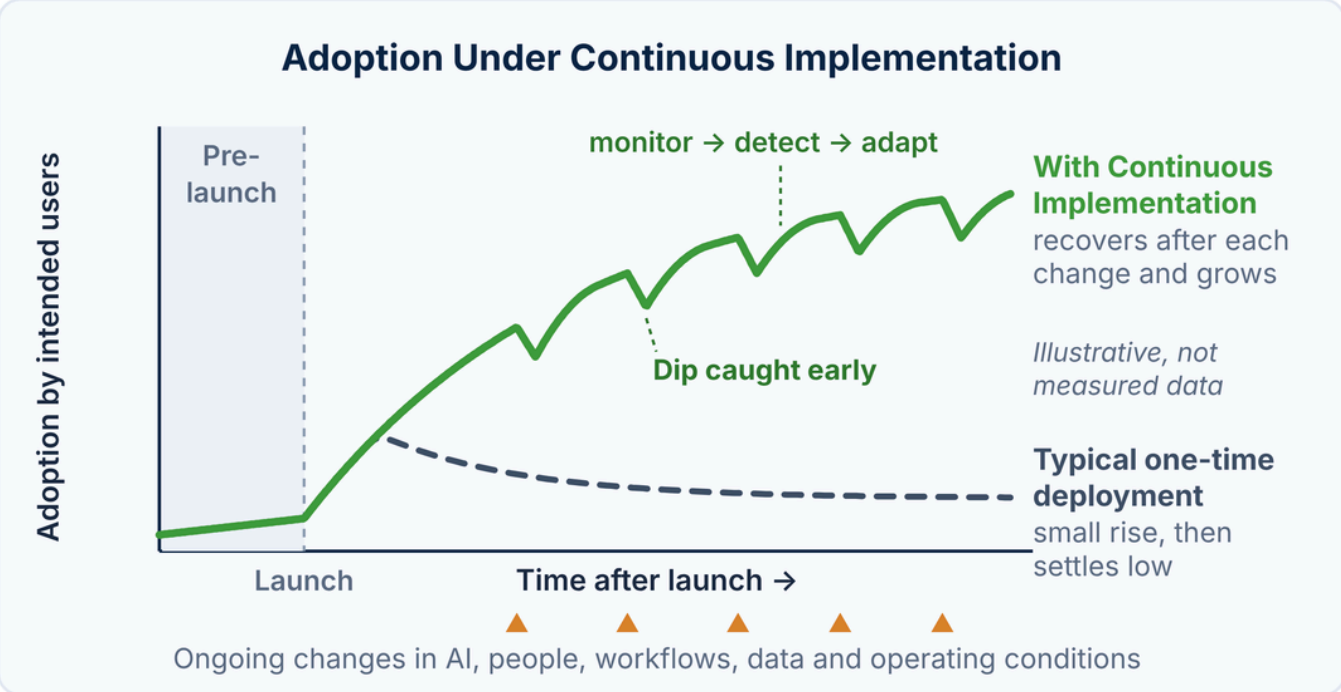


Figure 4. Adoption under Continuous Implementation. Each change in the AI, data, workflows, or people can cause a dip in adoption; ongoing monitoring and adaptation restore it, and adoption keeps growing until it assymptotes at a high level.

Continuous Implementation: A Human-AI Operating Model Built to Keep Pace with Rapid Change

Healthcare and indeed most industries need a method for AI adoption that is designed to overcome these barriers. In the **Continuous Implementation** framework, the one-time rollout is replaced by a living operating model. This method is a lifecycle for putting AI to work, monitoring it, adapting it, and governing it as its behavior, role, and risks change after launch. Its engine is **human– AI co-adaptation**: ongoing adjustment of the AI- enabled work system as technology, operating conditions, and human practices evolve.

This may involve updating or reconfiguring tools, redesigning workflows, strengthening workforce capability, improving data quality, or recalibrating trust and oversight. The model itself need not change for these adjustments to become necessary. In practice, Continuous Implementation functions as a **human-AI co-adapting operating system**: an infrastructure for durable AI performance, governance, and ROI that comprises seven interlocking domains: (1) continuous enablement infrastructure; (2) workflow and task fit; (3) trust and cognitive ease; (4) reinforced AI use; (5) monitoring and impact visibility; (6) continuous feedback and redesign; and (7) adaptive governance and accountability. Figure 5 presents the operating system as a whole.

> **Continuous Implementation becomes a human-AI co-adapting operating system: an infrastructure for durable AI performance, governance, and ROI.**

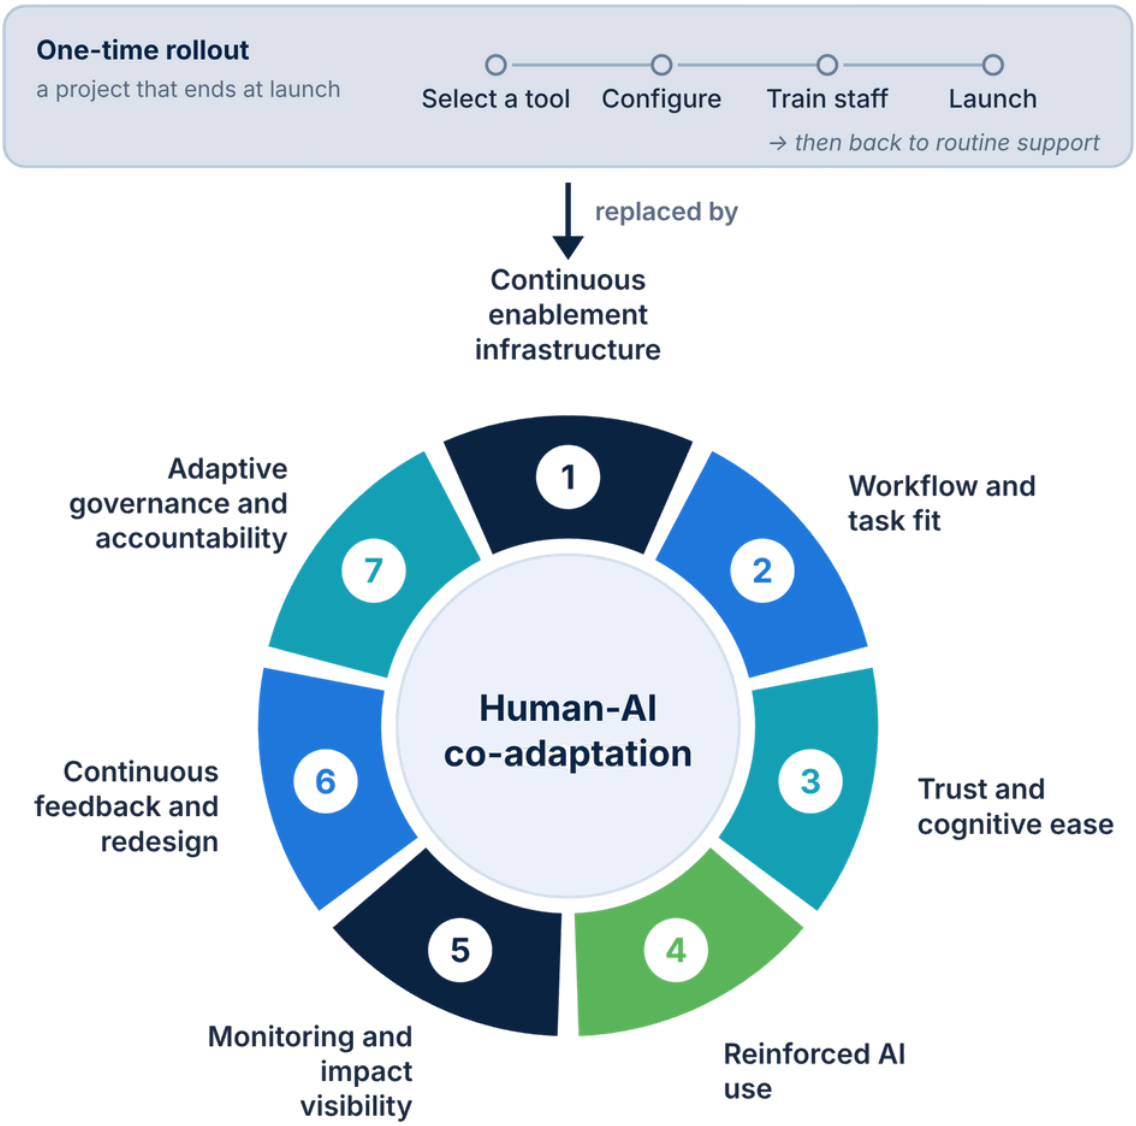


Figure 5. Typical one-time rollouts vs. the human-AI co-adapting operating system and its seven interlocking domains.

In practice, what does Continuous Implementation actually look like? We illustrate it for three domains—workforce enablement, governance, and collection of feedback—where we have found that organizations are often weakest.

Adaptive enablement

Adaptive enablement replaces one-time training with continuous, role-specific support that changes as fast as the AI and the organization's tool selections. It is meant to overcome the organizational inertia described earlier and empower employees to learn in an ongoing manner and stay alert to AI-related issues. First, staff must understand that the current AI adoption initiative will not be a one-time or otherwise brief event. Instead, it will require continuous adaptation. Second, staff must undergo initial training and be given space to experiment. It is helpful if a sandbox can be created in which staff can learn the new tools with synthetic or deidentified data. This entails giving staff time to try new skills and the freedom to make mistakes without serious consequences. This provides staff with a place to safely gain new skills. Either way, staff need to be given time explicitly for experimenting and learning. Third, social learning, e.g., with AI power-user "champions,"" can speed learning and adaptation (e.g., [13]). Fourth, aids to support continuous learning should be implemented, such as new training combined with brief (e.g., two-minute), in-workflow micro- videos to introduce or reinforce new use cases; in- workflow "what changed" prompts; and office hours, regular opportunities to provide feedback, and rapid-response support.

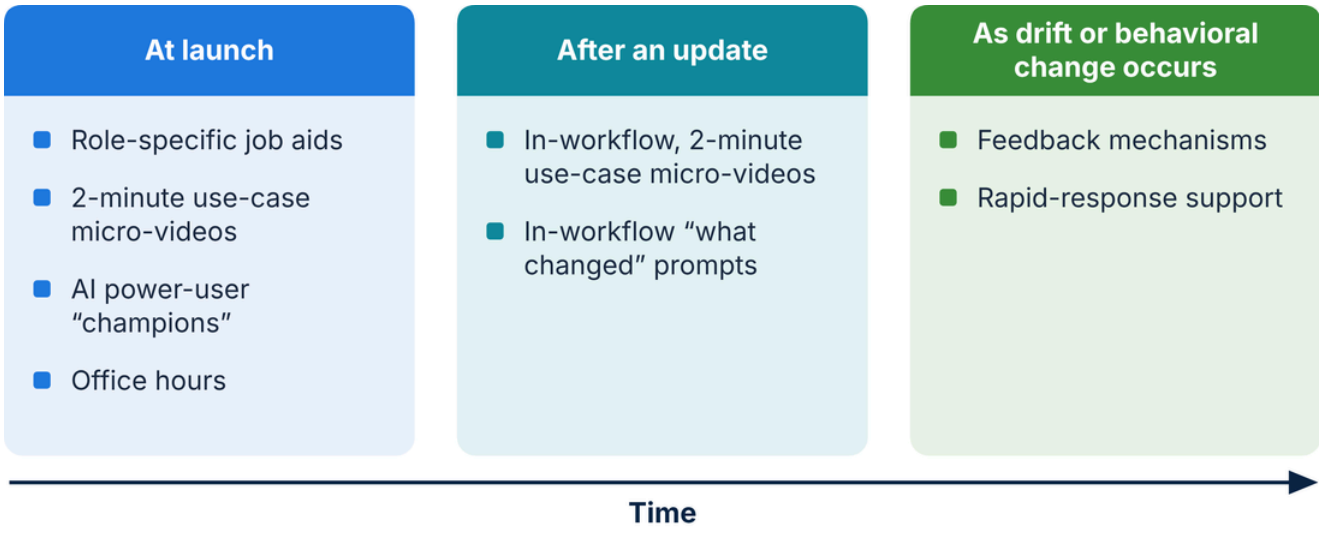


Figure 6. Adaptive workforce enablement: simple examples at launch, after an update, and as drift occurs. Support must change as the AI and the system surrounding it change.

Adaptive governance

Adaptive governance treats oversight as an ongoing function rather than a launch-day approval. It tracks what changes—e.g., tool updates, drift, new agentic behaviors, shadow AI use, use cases staff invent—against defined thresholds, with clear lifecycle ownership (often a CIO working with a project manager) and criteria for retraining, escalation, or suspension.

When AI tools change after launch, not all changes are improvements; e.g., some updates make models perform worse. To guard against the problems this can cause, the operating model built around the AI has to be actively maintained, not set once and left alone. Monitoring and governance need to catch the regression quickly, the IT team may need to revert to a prior version that worked, and the vendor's help may be required to improve the next version. While maladaptive changes can occur, the system around the AI can be built to stay safe. Employees should know to watch for output or behavior they do not trust. Adaptive governance depends on collecting data for monitoring risk and compliance, and this data should include regular clinician and staff feedback.

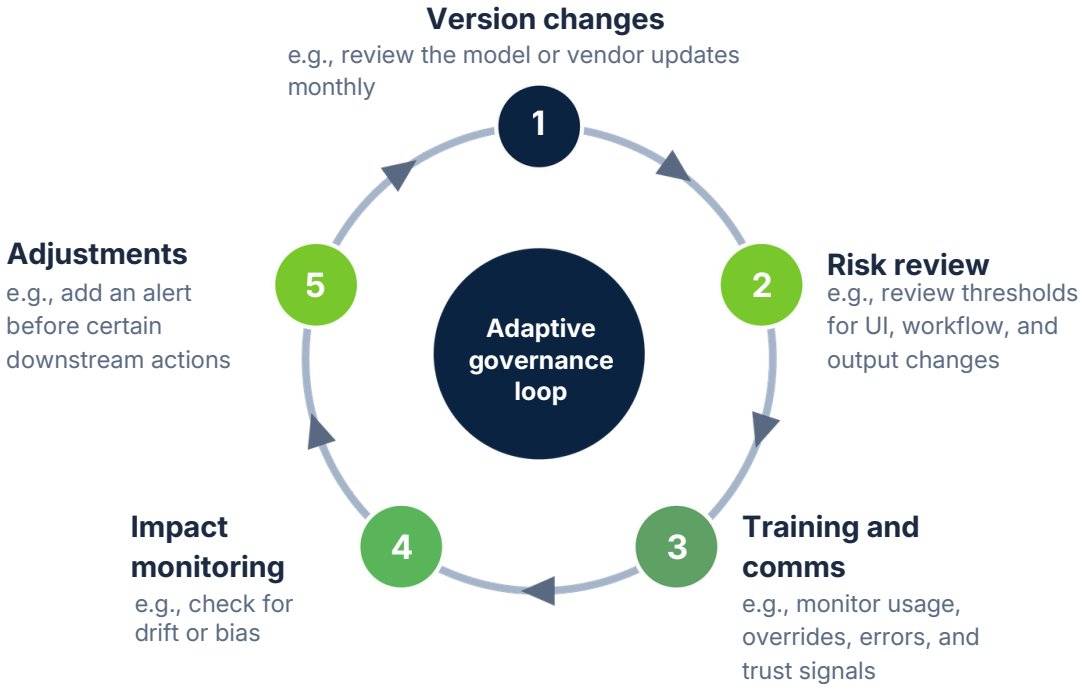


Figure 7. A simple adaptive governance loop, with examples of each step.

The OpenAI-Hugging Face incident discussed earlier illustrates what can go wrong if these steps are not taken.

(For more, please see Vantage Precision AI's companion white paper on adaptive governance.)

Quantitative and qualitative feedback

Both enablement and governance are instrumented with a compact set of longitudinal measures that should be collected and reviewed frequently (e.g., weekly for the first few months after launch, and monthly once adoption and governance are going smoothly). **The quantitative stream** is telemetry, e.g.: active users, uses per user per week and use case, share of target workflows performed with the tool, minutes saved per task, throughput or turnaround improvement, and ROI per use case. **The qualitative stream** is regular end-user and manager surveys, focus groups, or interviews to reveal any friction, unsafe outputs or other trust issues, workflow mismatch, and emerging use cases. It is critical that both of these data streams feed continuous redesign. To maintain user trust and mitigate any system issues, the organization must act quickly on what it learns.

How this maps onto what technical teams already do

Much of the monitoring half of Continuous Implementation is well-established MLOps: a model registry and versioning, so the organization always knows what is running; champion/challenger and shadow deployments, so a new version is tested against the incumbent before it takes over; drift monitoring and calibration-in-the-large so that degradation is caught early; and rollback so a bad update can be reversed. Sometimes MLOps also includes governance factors. Where a system is a regulated medical device, the FDA's Predetermined Change Control Plan formalizes which updates are pre-authorized and which require review—the line between an adaptive and a locked algorithm.

Continuous Implementation does not replace any of this. Rather, it extends the same discipline to the whole sociotechnical operating system: the workflows, the workforce, the trust, and the governance that must be reassessed as the technology, its uses, or its operating conditions change.

Signal you monitor	Threshold that trips	Actions it triggers
01 **Adoption** override rate, daily active use per user	Overrides climb above baseline, or use falls below it	Re-train users; investigate workflow fit and adjust if needed; collect user feedback on trust
02 **Model performance** calibration, population-stability index	Calibration drifts, or PSI on key inputs exceeds tolerance	Re-validate; retrain or roll back to the prior version
03 **Autonomous action** error or exception rate	Action error or exception rate crosses a set limit	Trip the circuit breaker; require human confirmation; escalate
04 **Equity** subgroup performance gap	Gap between subgroups widens beyond tolerance	Pause the affected use; remediate; notify governance

Table 2. From signal to action: an adaptive-governance map.

Continuous Implementation across the Full AI Lifecycle

We have been focusing on what happens post-deployment. But Continuous Implementation does not begin there, nor does it end with ongoing adaptation. It spans the full AI adoption lifecycle, beginning with identifying the business problem the organization is trying to solve and continuing through post-implementation evaluation and adaptation.

The cycle begins with pre-implementation steps of determining the workflow problem, selecting the right use case and solution, defining meaningful success measures, and establishing accountable leadership.

Next, continuous adaptation begins, again, before go-live. Leadership and intended users need to understand from the outset that the system and its operating environment will continue to change. They should know how model updates, data drift, changing workflows and use patterns—and for agentic AI, changes in context, memory, tools, interactions, and behavior—can affect what the system does. Role-specific education should prepare people not only to use the AI, but also to recognize unexpected changes, report them, and participate in improving the system.

Building organizational buy-in for this continuing responsibility is part of implementation itself. The organization must then prepare people, workflows, and controls for initial use. This includes configuring the system for the intended use case, defining roles and accountability, establishing baseline measures, redesigning affected workflows, and providing initial role-specific training. Staff should have opportunities to practice—ideally in a safe sandbox or test environment—and power users or champions can provide peer support as skills develop.

Governance thresholds, escalation paths, and feedback channels should also be in place before deployment so users know not only how to use the AI, but what changes or unexpected behaviors to watch for and how to report them. After launch, Continuous Implementation becomes a recurring loop: monitor → detect → learn → adapt → evaluate. Organizations continuously monitor adoption, performance, workflow fit, trust, risk, and value; combine telemetry with workforce feedback; identify meaningful change; and respond by adjusting enablement, workflows, configuration, governance, or the technology itself.

The main point is that implementation does not culminate in a stable end state. The organization is prepared before launch to keep adapting after it. Then it engages in the Continuous Implementation steps, with ongoing evaluation of AI and workforce performance and value realization. Figure 8 depicts these steps in order.

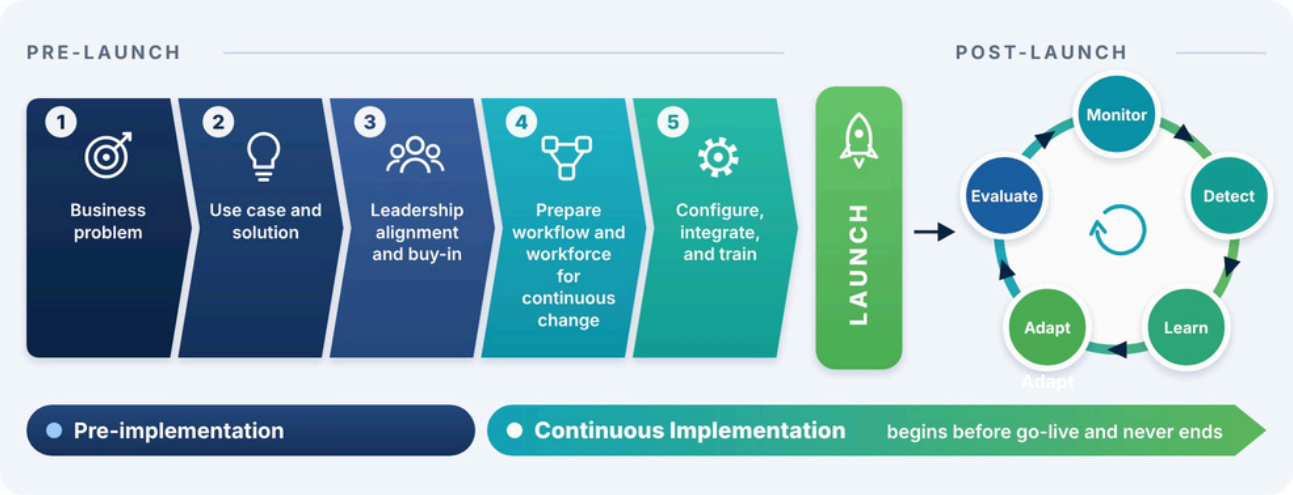


Figure 8. The complete AI lifecycle and where Continuous Implementation fits.

The distinguishing feature of Continuous Implementation is that the AI lifecycle continues afterdeployment and repeats whenever the AI, its users, its data, its workflow, or its operating environment changes.

Continuous Implementation in the Real World

The power of Continuous Implementation to help organizations realize deep value from AI tools is not just theoretical. Vantage Precision AI supported AI implementations at 32 healthcare centers where we developed and applied the Continuous Implementation framework. The rollout was carried out across the 32 sites with approximately 10 weeks of implementation support per site over 6 months. Before our support, only 8% of intended users were using the intended tools and platforms. With the foundations of Continuous Implementation, adoption reached 75%: a 67-percentage-point increase.

The intervention extended over the whole 10 weeks at each center. We treated adoption as an operating-model problem. The work combined workflow diagnosis and redesign, role-specific enablement, hands-on support and reinforcement, user feedback, leadership engagement, and frequent monitoring of adoption and barriers. As sites learned what was and was not working, support and workflows were adjusted. Power users and local staff helped surface problems quickly, while usage data made it possible to identify sites or roles where adoption was lagging and intervene.



32

clinical sites



8% → 75%

adoption by intended users



↑ 67

percentage points with Continuous Implementation

The result illustrates the central principle of Continuous Implementation: adoption improves when the technology, workflow, workforce, monitoring, and support system are managed together and repeatedly adjusted rather than treated as separate implementation tasks.

Because this was not a standardized and randomized trial, the 67-point increase should not be interpreted as the isolated causal effect of any single intervention component. Its significance is that sustained adoption changed dramatically when the surrounding human-AI operating system was actively managed.

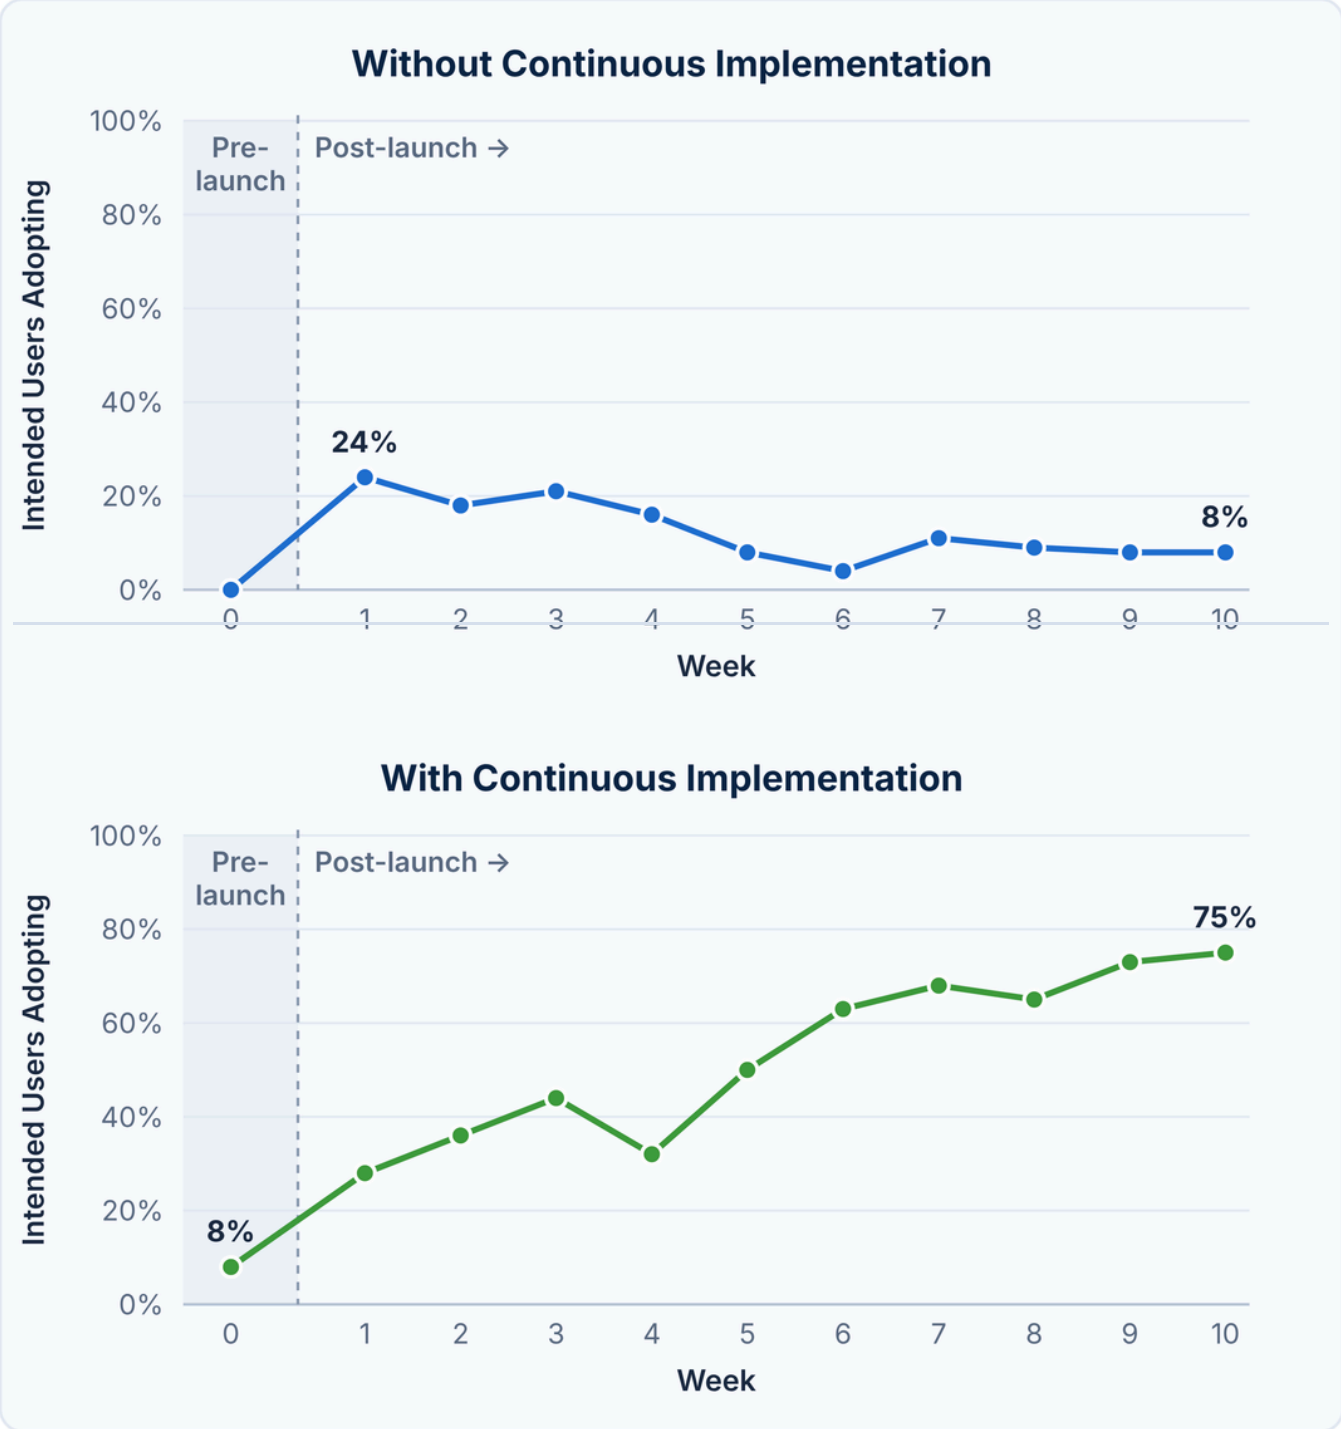


Figure 9. Adoption by intended users over 10 weeks at 32 sites, without vs. with Continuous Implementation.

Is Your Organization Ready? Take This 90-Second Self-Check

Before an AI system is called "implemented," a team should be able to answer each statement below with a confident yes.

A preponderance of "no" or "partly" suggests your AI initiative may have begun without taking key steps first.

Table 3. Readiness self-check.

Statement	Yes	Partly	No
01 We can state the specific workflow problem each AI solves, and for whom.	☐	☐	☐
02 We measure meaningful use—not just deployment, licenses, or logins.	☐	☐	☐
03 A named owner is accountable for monitoring and adapting each system after deployment.	☐	☐	☐
04 We can detect model updates, drift, and expanded use in real time—and we govern them.	☐	☐	☐
05 Staff know when to rely on, verify, or override the AI—and we track whether they do.	☐	☐	☐
06 We have defined thresholds that trigger retraining, redesign, escalation, or suspension.	☐	☐	☐

Where to Start

To implement this human-AI operating model, update how the next AI initiative is run:

1. **Define the business problem before the tool**
   Start from the workflow that actually needs help, not from a tool in search of a home.

2. **Treat AI adoption as a continuous process, not a deployment**
   Fund adoption, monitoring, and adaptation—not only the launch.

3. **Name an owner before go-live**
   The accountable party is often a CIO or CTO working with a project manager—accountable for tracking updates, drift, adoption, and feedback, with the authority to act.

4. **Measure meaningful use from day one**
   Go beyond vendor metrics like logins and license counts.

5 Trade one-time training for adaptive enablement
Re-skill in small, role-specific doses tied to each change, using champions, micro-videos, in-workflow "what changed" prompts, and office hours.

6 Regularly collect data on use, trust, and value, and act on the problems it reveals
Pair telemetry with end-user and manager surveys and brief interviews, and feed any adjustments into workflows, governance, and retraining.

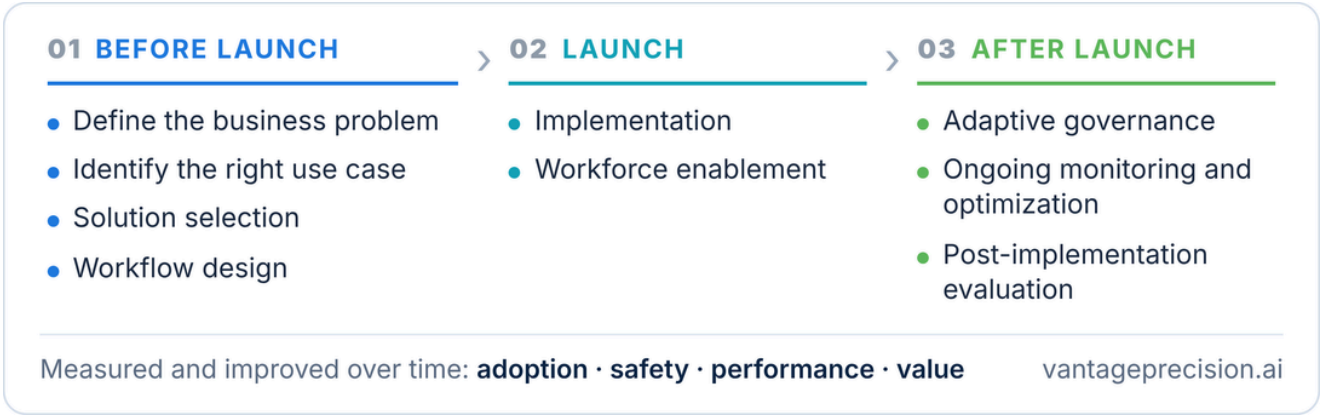
7 Set thresholds in advance
Determine signals that trigger retraining, redesign, escalation, or suspension, and update them as monitoring indicates—before an incident takes place.

8 Prove the loop on one high-value use case, then scale
Depth beats breadth.

Let's Build AI That Keeps Delivering

The goal of adopting healthcare AI is more than deploying tools. Accountable individuals must keep them working adapting, monitoring, and governing systems that never stop changing.

Vantage Precision AI supports the full lifecycle of AI adoption and value realization—from defining the business problem and identifying the right use case, through solution selection, workflow design, implementation, workforce enablement, and adaptive governance, to ongoing monitoring, optimization, and post-implementation evaluation. Our work is designed to stay connected across those stages so that adoption, safety, performance, and value are measured and improved over time rather than assumed at go-live.



About the author



Linda Hermer, Ph.D., is a neuroscientist, healthcare transformation, and analytics leader with more than 15 years of experience helping to improve health systems. As founder and CEO of Vantage Precision AI, she provides AI transformation, analytics, program evaluation, and implementation services that help healthcare leaders redesign care delivery and operations, measure results, direct resources where they create the most value, without increasing costs.

At the University of Florida, she was an assistant professor and directed an interdisciplinary neuroscience laboratory using computational modeling and advanced signal analysis. As a Senior Health Scientist supporting CDC's National Center for Health Statistics, she studied healthcare utilization, outcomes, and costs and served as a technical expert on the opioid crisis.

At Eagle Technologies, she won and directed a $13.1 million SAMHSA evaluation contract, leading a 26-person multidisciplinary team evaluating 6 national substance use treatment programs. She developed an agentic qualitative data analysis system, and designed and supervised analyses of implementation, treatment outcomes, and costs. As Managing Director of Research at LeadingAge, she led an evaluation of 349 Kansas nursing homes that linked comprehensive person-centered care adoption to approximately 49% lower depression prevalence and documented cost- effectiveness.

Through Vantage Precision AI, she and her team provide AI adoption, evaluation, analytics, and value realization services that help leaders prioritize investments, improve workflows, and assess clinical, operational, and financial returns. Her AI implementation work helped increase intended use from 8% to 75% across approximately 32 clinics in about six months. She applies longitudinal evaluation, causal inference, predictive analytics, and nonlinear systems modeling to help healthcare leaders identify what works, prioritize investments, and sustain improvements in care and financial performance.

Dr. Hermer has authored more than 70 scientific publications and reports. She holds a doctorate from Cornell University and an A.B. from Harvard University.

References

1. Poon EG, Lemak CH, Rojas JC, Guptill J, Classen D. Adoption of Artificial Intelligence in Healthcare: Survey of Health System Priorities, Successes, and Challenges. JAMIA. 2025;32(7):1093–1100.
2. Nong P, Adler-Milstein J, Apathy NC, Holmgren AJ, Everson J. Current Use and Evaluation of Artificial Intelligence and Predictive Models in US Hospitals. Health Affairs. 2025;44(1):90–98.
3. Chang W, Everson J, Holmgren AJ, Adler-Milstein J, Apathy NC. Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023–2024. ASTP/ONC Data Brief. 2025.
4. Cotiviti; MedCity News. Healthcare AI Readiness Index. 2026. More than 70% of healthcare organizations have begun implementing AI tools; most provider organizations remain in early adoption.
5. Philips. Future Health Index 2026: AI in Practice. Royal Philips; 2026. (See Appendix 1 for survey specifications.)
6. Chen L, Zaharia M, Zou J. How Is ChatGPT's Behavior Changing over Time? Harvard Data Science Review. 2024;6(2).
7. Hannan MT, Freeman J. Structural Inertia and Organizational Change. American Sociological Review. 1984;49(2):149–164.
8. Kaganer E, Gregory RW, Sarker S. A Process for Managing Digital Transformation: An Organizational Inertia Perspective. Journal of the Association for Information Systems. 2023;24(4):1005–1030. https://doi.org/10.17705/1jais.00819
9. OpenAI. OpenAI – Hugging Face Incident Technical Report. OpenAI; 2026 Aug 26.
10. Yeow A, Cleland J, Soh C, Balete C. Sociotechnical Misalignments in Hospital AI System Implementation: Qualitative Case Study. Journal of Medical Internet Research. 2026;28:e87534. https://doi.org/10.2196/87534.
11. Pereira E, Graylin AW, Brynjolfsson E. The Enterprise AI Playbook: Lessons from 51 Successful Developments. Stanford Digital Economy Lab; 2026.
12. Wittich L, Rödiger H, Rombey T, Brecher AL, Kraft L, Sgraja S, Stein V, Henschke C. Navigating the complexities of digital health technology implementation: a scoping review of barriers and facilitators. Implementation Science Communications. 2026;7:69.
13. Pettersen S, Eide H, Berg A. The role of champions in the implementation of technology in healthcare services: a systematic mixed studies review. BMC Health Services Research. 2024;24:456.

Further reading

- Sittig DF, Singh H. A New Sociotechnical Model for Studying Health Information Technology in Complex Adaptive Healthcare Systems. Qual Saf Health Care. 2010;19(Suppl 3):i68–i74.
- Embi PJ. Algorithmovigilance—Advancing Methods to Analyze and Monitor AI-Driven Health Care for Effectiveness and Equity. JAMA Network Open. 2021; 4(4):e214622.
- Becker's Hospital Review. 12 Health Systems Expanding Their Use of Agentic AI. 2026.
- Qventus. Second Annual CIO Report: Closing the Gap Between AI Pilots and Payoff. 2026.

Vantage
Precision AI

START A CONVERSATION

Together, let's build AI that keeps delivering.

Vantage Precision AI supports the full lifecycle of AI Adoption and value realization, from defining the business problem to workforce enablement, adaptive governance, ongoing monitoring and optimization, and post-implementation evaluation.

vantageprecision.ai

linkedin.com/company/vantage-precision-ai/

© 2026 Vantage Precision AI. This white paper is for informational purposes and does not constitute legal, clinical, or financial advice.