Massachusetts League of Community Health Centers

AI Strategy-to-Implementation Playbook

Empowered Health Centers. Collective Action. AI That Works for CHCs.

Prepared for the Massachusetts League of Community Health Centers

Prepared by Phicil-itate Change LLC

Built on SEEDS of Innovation™

Version 1.0 · July 2026

Disclaimer

This Playbook is a general governance resource for Massachusetts community health centers. It is a compilation of practical guidance, templates, and decision tools — not legal, clinical, financial, or compliance advice. Nothing in this Playbook should be read as Massachusetts League of Community Health Centers (MLCHC) or Phicil-itate Change LLC affirmatively endorsing any specific artificial intelligence (AI) use, product, practice, or policy as compliant with applicable law. Each health center is solely responsible, in conjunction with its own independent legal counsel, for determining whether its AI uses and policies — including vendor contracts, executed Business Associate Agreements, written information security program, data minimization, technical safeguards for encrypting data both at rest and in transit, restrictions on AI model training using Personal Information, prohibition of autonomous AI decision making without medical professional review, clinical documentation practices, and claims submission — comply with applicable state and federal law. AI use does not change a health center's underlying legal and compliance obligations; AI-assisted decisions, documentation, and claims are held to the same standards as any other.

AI is a fast-developing technology, and regulatory and policy changes are expected to be implemented in the coming years to catch up. This Playbook reflects best practices and guidance available as of its publication date and may not be current depending on when a health center reviews it. Health centers should always ensure they are reviewing the most up-to-date version of the Playbook and confer with legal counsel to determine how any recent or anticipated regulatory or policy changes could affect them.

Use of this Playbook does not guarantee compliance with any law, regulation, contractual obligation, or payer requirement. MLCHC and Phicil-itate Change LLC are not liable for outcomes arising from a health center's use of this Playbook. Any vendors or products referenced in case examples are included for illustrative purposes only and do not constitute an endorsement. The tools, templates, and frameworks provided directly in this Playbook are intended for health centers to use and adapt based on their own circumstances.

This Playbook was developed through document review, stakeholder interviews, and an environmental scan of Massachusetts community health centers. It does not include direct research with patients; where patient perspective is referenced, it reflects information gathered from health center staff, leadership, and existing organizational sources.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Copyright and Trademark

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC. All rights reserved.

SEEDS of Innovation™ is a trademark and proprietary methodology of Phicil-itate Change LLC, separate from the copyright in this publication. No act of accessing, downloading, reproducing, or sharing this publication, and no permission granted to do so, transfers, licenses, or otherwise grants any right to use SEEDS of Innovation as a trademark, service mark, certification mark, or branding element. Use of SEEDS of Innovation as a service name, publication title, certification mark, product identifier, platform name, training brand, or commercial identifier requires prior written consent from Phicil-itate Change LLC. Reference to SEEDS of Innovation in published materials does not imply endorsement, sponsorship, or affiliation by Phicil-itate Change LLC.

How to Cite This Playbook

Massachusetts League of Community Health Centers. MLCHC AI Strategy-to-Implementation Playbook. Prepared by Phicil-itate Change LLC. 2026.

Use the ™ symbol on first reference to SEEDS of Innovation and attribute ownership to Phicil-itate Change LLC. Accurate, attributed reference in academic, journalistic, organizational, or governmental materials is permitted. Use in an organization's own publications, service descriptions, training offerings, consulting materials, or software requires prior written permission from Phicil-itate Change LLC.

Version Control

Version	Date	Section(s) Changed	Summary of change(s)	Author of Changes
v1.0	July 10, 2026	All	Initial publication	Phicil-itate Change LLC

Acknowledgments

This Playbook was developed by Phicil-itate Change LLC for the Massachusetts League of Community Health Centers. Each contributor is recognized individually below.

Sheila Phicil, MPH, MS, PMP, FACHE Project Lead and Senior Implementation Advisor. Creator of SEEDS of Innovation™ framework and led the overall direction, stakeholder engagement, and final integration of the Playbook.

Srushti Madhure, MS Health Informatics Specialist. Led the technical infrastructure, data readiness, and vendor-evaluation components, and the HIPAA and Massachusetts privacy alignment. Designed visual resources and ensured they followed accessibility best practices for clear and inclusive communication.

Natalie Johnson, MPH Project Manager and Implementation Lead. Managed the workplan, timelines, and trainings; supported coordination of stakeholders across the Mass League and Phicil-itate Change teams.

Sarah Faulkner, MBA, PMP Implementation Advisor and Scrum Lead. Led Agile delivery and shaped the playbook's implementation guidance, change-management tools, and measurement approach.

Rosa M. Colón-Kolacko, PhD, MBA, CDM, SHRM-SCP Senior Consultant for workforce development, change management, health equity and Human-Centric AI design. Provided foundational guidance on AI leadership readiness, change management, and AI healthcare impact assessment.

With thanks to the Massachusetts League of Community Health Centers, including Susan Adams, Linda Naval, Tinamarie Fioroni, and Tarsha Weaver, and to the health center leaders and frontline staff whose input greatly shaped the content, structure, and design of this playbook.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Table of Contents

Contents

Disclaimer 2
Copyright and Trademark 3
How to Cite This Playbook 3
Version Control 3
Acknowledgments 4
Table of Contents 5
Opening Letter 8
Executive Summary 9
Why this Playbook, and why now 9
How the Playbook decides: SEEDS of Innovation™ 9
SEEDS connects strategy to action 10
The four phases 10
The AI Governance Hub: where this all runs 12
This is collective action across the network 12
Section 1: How This Playbook Works 14
1.1 AI Strategy and Workforce Readiness 14
1.2 How to Navigate This Playbook 15
1.3 The Decision Method: SEEDS of Innovation™ 17
1.4 The Implementation Approach 18
Section 2: Getting Started 20
2.1 Current-State Assessment 20
2.2 SEEDS Decision Reasoning Maturity 21
2.3 Minimum Viable Governance 22
2.4 The AI Governance Hub 23
2.5 90-Day Starter Roadmap 25
Approved-Use Inventory Checklist 29
AI Policy Checklist 34
AI Governance Hub Getting Started Kit 41
Leadership Talking Points Template 48
Section 3: Phase 1 — Listen and Learn 56
3.1 What “Listen and Learn” Means 56
3.2 Signal Types 57

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

3.3 Patient, Staff, and Board Voice 58
3.4 Monitoring Regulation, Vendors, and Payers 59
AI Signal Intake Form Template 62
Patient AI Feedback Form Template 67
Regulatory and Compliance Monitoring Checklist 71
Section 4: Phase 2 — Assess and Decide 77
4.1 Vendor Evaluation 77
4.2 Applying SEEDS to Signals 78
4.3 Maturity Check 80
4.4 Behavioral Health, SDoH, and Sensitive Populations 80
4.5 Pre-Deployment Impact Checklist 81
Vendor Evaluation Checklist 85
AI Signal Decision Assessment 98
Pre-Deployment Impact Checklist 109
Section 5: Phase 3 — Plan and Implement 118
5.1 From Decision to Designing a Plan 118
5.2 Components of a Good Plan 119
5.3 Workforce Lens and Role Changes 120
5.4 Communication with Staff and Patients 120
5.5 Implementation Readiness 121
Design Plan Template 124
AI Implementation Readiness Checklist 136
Section 6: Phase 4 — Steward 141
6.1 Measuring Performance and Impact 141
6.2 Adjust, Scale, or Stop 142
6.3 Feeding Stewardship Back to the AI Governance Hub 143
Stewardship Review Checklist 146
Part A — Performance Measures 147
Part B — Review and Decision 150
Case Example Sharing Template 155
Section 7: Keeping the System Running 162
7.1 Regulatory Awareness and Proactive Monitoring 162
7.2 Keeping the Playbook Current 163
7.3 Network Feedback and Shared Learning 163
Appendix A: SEEDS of Innovation™ 165

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

What SEEDS Is 165
The Five SEEDS Pillars 166
Reasoning Maturity: How Deeply You Listen 167
Learn More 168
Appendix B: Implementation Approach Reference 169
Part 1: The CHC Tailoring Methodology 169
Part 2: Framework-by-Framework Reference 171
Supporting Frameworks 177
Complementary Framework 180
Appendix C: Glossary 181
Appendix D: Tools, Templates and Resources 189
Part 1: Playbook Tool and Template Index 189
Part 2: Resources for AI in Healthcare 189
Appendix E: AI Literacy Resources 196
Tier 1: Built for Community Health Workers and Healthcare Staff 196
Tier 2: Nonprofit and Mission-Driven 198
Tier 3: Universal and Foundational 198
Tier 4: Emerging (Watch List for MLCHC) 200
Summary at a Glance 201
Recommended Deployment Strategy 201

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Opening Letter

Dear Community Health Center leaders,

For more than 60 years, community health centers have met our nation's greatest challenges with innovation, determination, and an unwavering commitment to community. Today, we face another such moment. The advent of Artificial Intelligence (AI) offers real promise to meet many of the challenges faced by health centers, while at the same time raises important questions about governance, trust, equity, impact, and return on investment.

You have asked for the practical tools, templates, and guidance to implement AI with both urgency and intent. Importantly, you have highlighted the need for change management, education, and a formalized strategy driven by stated goals, measurable objectives and organizational alignment to ensure new technologies empower and support the people who deliver care. That is why we've developed the **Mass League AI Playbook.**

The AI Playbook is a vital component of our collective response to innovation and a key mechanism for operationalizing the Mass League's broader ***Collective Action Strategy to Shape AI for CHCs*.** Rooted in evidence and extensive engagement with health center members, it translates strategy into action by helping health centers evaluate, adopt, and govern AI in ways that are grounded in workforce education, patient trust, and a clear commitment to putting mission, people, and community first.

AI cannot exist as an isolated IT initiative; it must be intentionally integrated across an organization. To further that integration, this AI Playbook was developed with two distinct purposes: to help individual health centers build the internal knowledge, strategic road map and governance needed to implement AI responsibly, unique to their organization; and to help health centers act in unison to shape AI solutions that work specifically for CHCs, including their workforce and patients.

That work is significant, and this AI Playbook is equally comprehensive. At over 200 pages, it serves as an extensive resource, providing a full suite of tools, templates, and governance structures designed to support your team's efforts, for every stage of a health center's AI adoption journey. Paired with live train-the-trainer sessions and implementation support, we are committed to working alongside you to ensure that this is a living, breathing document that is worth your time and investment.

We extend our deep gratitude to Mass General Brigham for making this Playbook possible and for sharing lessons from its work to advance responsible AI in care delivery. Their partnership reflects the ecosystem collaboration needed to ensure community health centers are not left behind in this digital evolution. We would also like to thank Phicil-itate Change for their collaboration and expert guidance in shaping the development of our Playbook.

As AI evolves, our strategic approach will evolve too. We will continue to listen, learn, and support health centers through workforce adoption, shared learning, partnerships, advocacy, and purchasing opportunities. Together, we will not simply adopt AI. We will shape it.

Sincerely,

Michael Curry, Esq.

President & CEO, Massachusetts League of Community Health Centers

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Executive Summary

MLCHC AI Strategy-to-Implementation Playbook

Artificial intelligence (AI) is already in your health center. It arrives through vendor updates, through staff trying new tools, and through features already built into the software you use every day. The question facing health center leaders is no longer whether to engage with AI. It is whether each AI decision gets made with intention, evidence, and care for the people it affects.

This Playbook gives your health center a way to do that. It turns AI from something that happens to your organization into something your organization governs. It is the shared roadmap behind MLCHC's AI strategy, built so health centers can act together.

Why this Playbook, and why now

Health centers are being asked to do more with the same staff, the same time, and the same resources. AI tools promise to help. Some will. Others will add work, create risk, or quietly drift away from the reason you adopted them. Telling the difference takes a clear process, applied the same way every time.

MLCHC's AI strategy rests on three principles: Mission, People, and Community First; Collaborate and Align; and Lead with Urgency and Intent. This Playbook puts those principles into daily practice. It works whether you are exploring your first AI tool, reviewing one that arrived without a formal decision, or checking whether a tool already in use still earns its place.

This Playbook is grounded in evidence. Before any tool was written, Phicil-itate Change conducted an environmental scan of what Massachusetts health centers are experiencing with AI. It drew on member surveys, input from frontline staff and health center leaders, and a wide review of the policy, vendor, and regulatory forces shaping the field. The scan confirmed what MLCHC's strategy already set out to meet: health centers are adopting AI quickly, and they have a real opportunity to lead by building strong, shared governance as they go. This Playbook turns that opportunity into a practical path, helping each health center put governance in place before the next AI decision arrives, and building patient voice into AI decisions from the start. The full environmental scan is available from MLCHC on request, as shared context for health centers putting this Playbook to work.

How the Playbook decides: SEEDS of Innovation™

Every tool in this Playbook runs through one decision standard, called SEEDS of Innovation™. Developed by Phicil-itate Change, SEEDS is built on five pillars, each a question your health center answers, in order, before it builds, buys, or keeps using any AI tool:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- **Systemic Signal Assessment**: Understand the system you are trying to change, and listen to the people closest to it.
- **Everyone-Centered Design**: Design for the people who face the most barriers, so the tool works for everyone.
- **Ecosystem Alignment**: Make sure the tool fits the whole system it joins.
- **Data Sovereignty**: Keep ownership of your data, and a share of the value it creates, close to the people whose work and lives created it.
- **Stewardship**: Stay responsible for the tool long after it goes live, and hold it to the reason you adopted it.

A tool can be technically strong, recommended by peers, and backed by its vendor (the company providing it), and still fail once it meets the real conditions of a community health center. SEEDS measures whether the system you actually run can carry the tool you are considering. When you answer these five questions before you commit, every AI decision rests on evidence instead of hope.

That is why SEEDS leads where health centers are already working to go. MLCHC's strategy names three goals for AI: protect access to quality care, restore time for care, and strengthen financial sustainability. SEEDS keeps every AI decision pointed at those goals, and at the people every health center serves.

SEEDS connects strategy to action

Your strategy sets what your health center wants AI to accomplish. SEEDS translates that into a sound decision on each tool, and then into an implementation your health center can stand behind. Underneath that decision standard, the Playbook draws on widely recognized public frameworks for AI governance in healthcare, the same thinking larger health systems are using. These include the National Institute of Standards and Technology (NIST) AI Risk Management Framework, the Coalition for Health AI (CHAI), the Health Sector Coordinating Council (HSCC) guidance on third-party AI risk, and the patient safety standards of ECRI (formerly the Emergency Care Research Institute), along with others that cover quality and readiness. Those frameworks were built for organizations with more staff and resources than most community health centers have, so the Playbook adapts each part to fit what a lean health center team can actually do, while keeping the protections that matter most for the patients health centers serve. The result keeps your health center current with the best thinking in the field, in a form built for your setting. The full set of frameworks behind each tool is listed in the Playbook's appendices.

The four phases

The Playbook moves every AI decision through four phases. Each phase begins with an AI signal, which is any indication, from a person, a system, a vendor, a regulator, or a peer organization, that something AI-related needs attention. The figure below shows how the tools carry a signal from one phase to the next.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

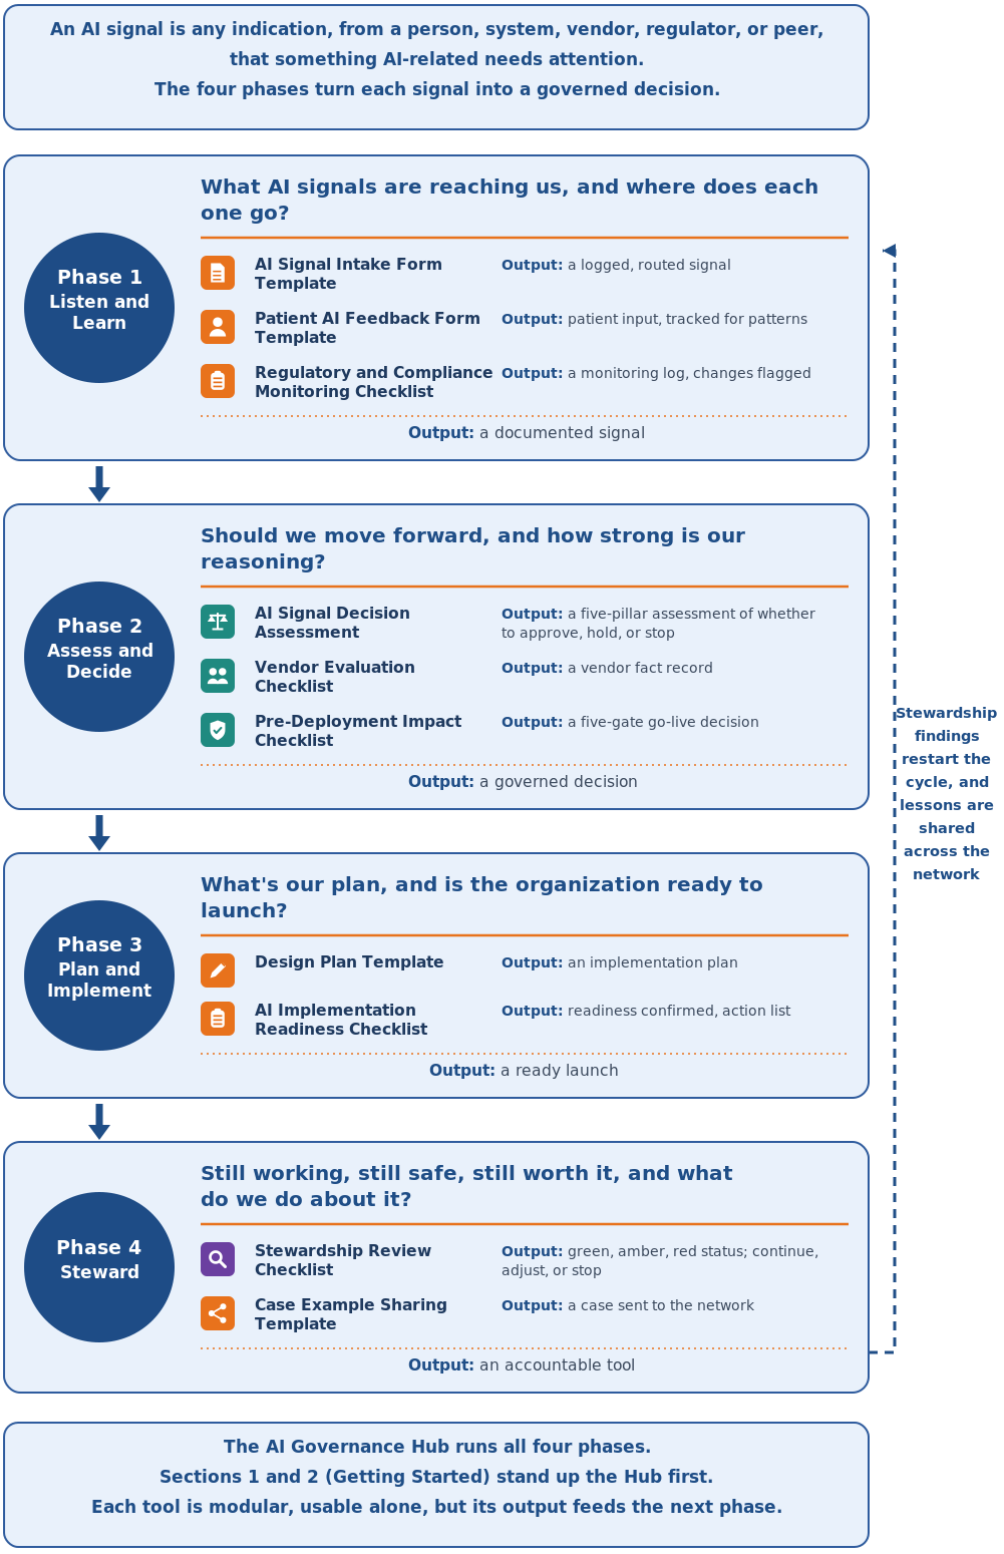

Figure ES.1. The Playbook at a Glance: Four Phases and the Tools That Move Each Decision Forward

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

In **Listen and Learn**, your health center catches these signals from staff, patients, vendors, and regulators, and routes each one to the right place. In **Assess and Decide**, your team weighs the evidence through SEEDS and reaches a clear recommendation: approve the tool (outright or with conditions), hold it for more information, or do not proceed. In **Plan and Implement**, an approved decision becomes a real plan, and your organization confirms it is ready to launch. In **Steward**, your team watches whether each tool already in use still works, still keeps patients safe, and is still worth it, then decides whether to continue, adjust, or stop it.

The AI Governance Hub: where this all runs

One structure makes the whole Playbook work: the AI Governance Hub. It is a standing team inside your health center, and it is where AI signals get reviewed, decisions get made, and outcomes get tracked over time. Every AI decision moves through it.

The AI Governance Hub brings together people from across the health center, because AI touches clinical care, operations, workforce, finance, compliance, and patients. Putting those perspectives in one room is what lets a health center see the whole picture before it commits. Two people keep the AI Governance Hub running: a Hub Lead, who owns the agenda and keeps decisions moving, and a Hub Coordinator, who handles scheduling and documentation. In a smaller health center, one person can carry both roles. The Hub runs on a steady rhythm: a short weekly check on new signals and open decisions, and a longer monthly review of the tools already in use.

Standing up the AI Governance Hub is the first thing your health center does, and 1 and Section 2 walk you through it. Once it is in place, the four phases give it a repeatable way to turn any AI signal into a sound, documented decision. That is how a health center moves from reacting to AI to governing it: one team, one shared standard, applied the same way every time.

Sections 2 through 6 each end with a hypothetical case example that follows the same health center using that section's tools from start to finish. If you want to understand a tool quickly, start there. The case examples walk through the tools in plain terms, decision by decision.

This is collective action across the network

No health center faces these decisions alone. When one health center learns something, through a tool that worked, a vendor that failed, or a decision to wait, that lesson moves through MLCHC to every other health center in the network. The Playbook builds that sharing into how it works.

Shared standards, applied across many health centers, also give the field strength that no single center has on its own. When every center asks an AI vendor the same protective questions, those expectations become harder for any vendor to work around. This is how the network moves from collective action to

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

collective power: shared purchasing strength, careful stewardship of patient data, and fair recognition for the knowledge and work health centers bring to the AI tools they help shape.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 1: How This Playbook Works

GOAL OF THIS SECTION

Find where to start, and understand the decision standard and implementation approach behind every tool in this Playbook.

WHAT'S COVERED IN THIS SECTION

1.1 AI Strategy and Workforce Readiness
Ground AI decisions in strategy and readiness

1.2 How to Navigate This Playbook
Find where to start for your situation

1.3 The Decision Method: SEEDS of Innovation™
Apply the five-question decision standard

1.4 The Implementation Approach
See the evidence base behind every tool

RELATED APPENDICES

Appendix A — SEEDS of Innovation™
The five pillars and maturity levels.

Appendix B — Implementation Approach Reference
Frameworks and eight filters behind the tools.

Appendix C — Glossary
Plain-language definitions of key terms.

Appendix D — Tools, Templates and Resources
Full index of every tool, in use-order.

Appendix E — AI Literacy Resources
Learning materials for leadership, clinical, and operations teams.

1.1 AI Strategy and Workforce Readiness

Every health center operates with a strategy: a set of commitments about who it serves, what quality of care it aims to provide, and what kind of organization it wants to be. AI decisions work best when they connect to that strategy. When a health center can point to its own goals as the reason it is evaluating a tool, governance becomes easier to explain, easier to sustain, and more likely to produce results that serve the center's mission, its people, and its community.

IN PLAIN TERMS

Artificial intelligence, or AI, is software that uses patterns in data to make suggestions or decisions. An AI tool might read a clinical note, flag a billing concern, schedule appointments, or sort through options to support a staff member's judgment. People build and train AI systems, and the choices made in that process shape what a tool does well and where it falls short. Full definitions of key terms appear in Appendix C: Glossary.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

A working knowledge of AI across the health center's teams strengthens that foundation. Leaders, clinical staff, and operations teams who understand how AI works, where it tends to produce errors, and how vendors represent their tools are ready to ask sharper questions, recognize misalignment early, and take part in governance with confidence. Building this knowledge is itself a form of collective action: health centers that develop AI readiness together negotiate from a stronger shared position.

Resources to support that learning are in Appendix E: AI Literacy Resources. They include materials for leadership, clinical teams, and operations staff, along with frameworks for understanding AI risk and vendor accountability in health center settings.

1.2 How to Navigate This Playbook

AI is already entering health centers across Massachusetts and beyond, through vendor proposals, staff experiments, and features built into software your organization already uses. This Playbook meets your health center wherever it stands today.

The Playbook is organized around four phases: Listen and Learn, Assess and Decide, Plan and Implement, and Steward. Together they cover the full life of any AI decision, from the first question about whether a tool fits your health center through ongoing oversight once it is in use.

Every health center begins by standing up the Section 2 — Getting Started foundation, which puts basic governance in place. From there, you focus on the phase that fits your situation. The figure below shows the whole Playbook at a glance.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

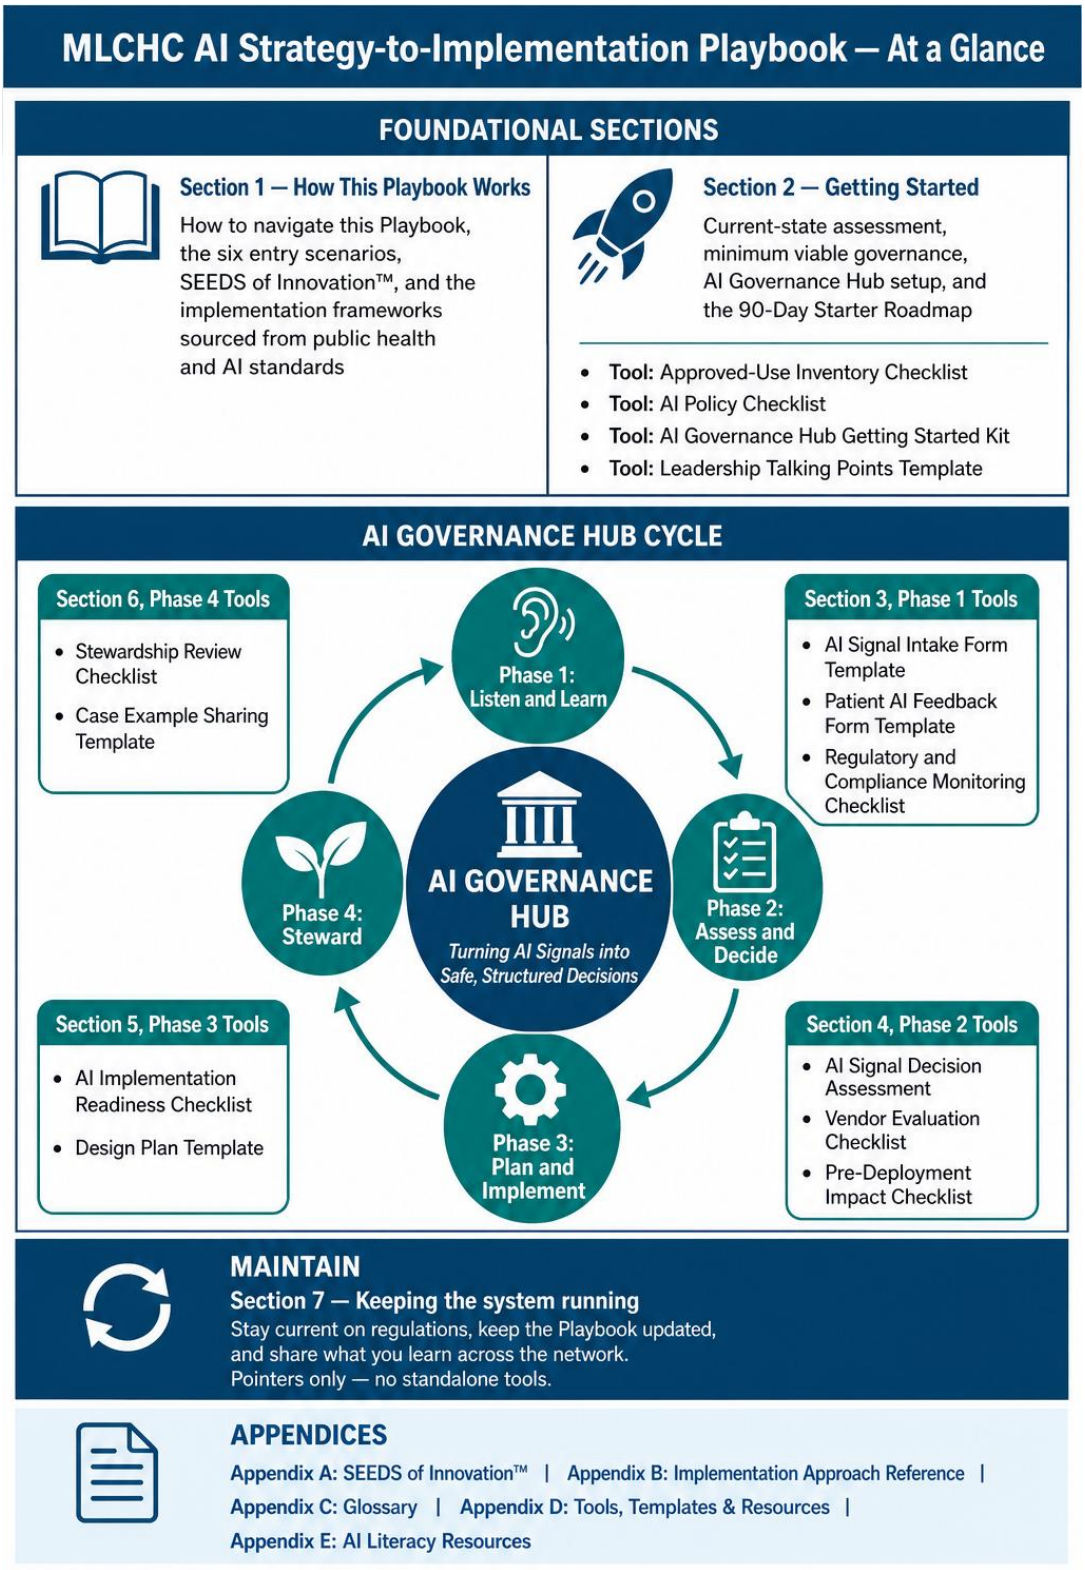

Figure 1.1. The Playbook at a Glance

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Use the entry scenario that matches where your health center is right now:

- **Exploring whether an AI tool is the right fit?** Begin at Phase 1 — Listen and Learn (Section 3).
- **Evaluating a specific AI tool?** Start at Phase 1 — Listen and Learn (Section 3), then move to Phase 2 — Assess and Decide (Section 4).
- **Formalizing an existing AI tool?** Start at Phase 1 — Listen and Learn (Section 3) and carry it into Phase 2 — Assess and Decide (Section 4).
- **Reassessing a changing AI tool?** Go to Phase 2 — Assess and Decide (Section 4).
- **Checking an approved AI tool?** Go to Phase 4 — Steward (Section 6).
- **Replacing or exiting an AI tool?** Go to Phase 4 — Steward (Section 6) for the exit. If you are replacing the tool, also start the incoming tool at Phase 2 — Assess and Decide (Section 4).

A full index of every tool, organized by section and phase, is in Appendix D: Tools, Templates and Resources.

1.3 The Decision Method: SEEDS of Innovation™

Every tool in this Playbook runs through the same decision standard: SEEDS of Innovation™, developed by Phicil-itate Change. SEEDS gives your AI Governance Hub a shared set of questions to answer before your health center builds, buys, or continues using any AI tool.

The connection to your work is direct. A tool can be technically strong and well-supported by its vendor and still fall short once it meets a community health center's real conditions. SEEDS tests whether the system your health center actually operates can carry the tool. When that test happens before you commit, every AI decision rests on evidence rather than assumption, and the payoff shows up where it matters: more time for frontline staff, better care for patients, and a stronger financial footing for the health center. Used consistently across the field, SEEDS also advances the shared strategy Massachusetts health centers are building together.

SEEDS works through five questions, answered in sequence.

- **Pillar 1: Systemic Signal Assessment.** Understand the system you are trying to change, and listen to the people closest to it, before acting.
- **Pillar 2: Everyone-Centered Design.** Design for the people who face the highest barriers, so the solution works for everyone.
- **Pillar 3: Ecosystem Alignment.** Make sure the solution fits the whole living system it joins.
- **Pillar 4: Data Sovereignty.** Keep ownership of data, and a share of the value it creates, close to the people whose lives and work created it.
- **Pillar 5: Stewardship.** Stay responsible for the solution long after launch, and hold it to the reason you built it.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

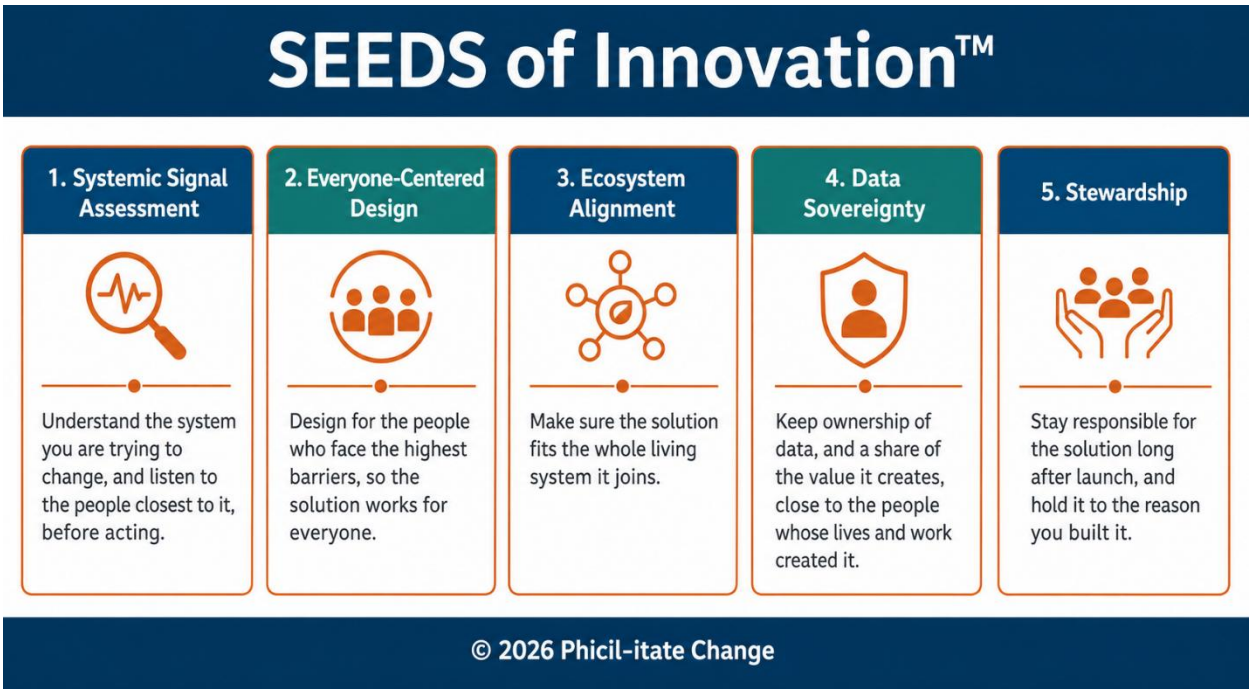

Figure 1.2. SEEDS of Innovation™: Five Pillars

Appendix A: SEEDS of Innovation™ gives the full pillar detail and the maturity levels that show where your health center stands within each pillar. Section 2 — Getting Started puts those maturity levels to work, and Section 4 — Assess and Decide puts all five pillars to work through the AI Signal Decision Assessment for every tool your health center evaluates.

1.4 The Implementation Approach

The tools in this Playbook draw on public frameworks for AI governance in healthcare, mapped in Appendix B: Implementation Approach Reference.

These frameworks were built for health systems with more resources than most community health centers. So the team preparing this Playbook did that adaptation work on behalf of health centers: it applied eight filters to every framework element before including it in a tool. Where a framework step assumed a dedicated role, such as an AI ethics officer, the Feasibility filter triggered an adaptation sized for lean health center teams. Each tool in the Playbook already accounts for all eight filters, so it meets the needs of a community health center as it comes.

- **Feasibility.** Can a lean health center team actually do this?
- **Harm and Stop-Gate Logic.** Does this protect the populations community health centers serve?
- **Burden and Sovereignty.** Does it close a care loop and preserve health center control?
- **Maturity Calibration.** Does this match the governance capacity most community health centers have?

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- **Scar Tissue Sensitivity.** Does it resemble a pattern health centers have faced before?
- **Evidence Legitimacy.** Was this developed for or tested in safety-net settings?
- **Labor and Workforce Trust.** Does this respect frontline staff as decision-makers?
- **Collective Action Value.** Does it protect something the field cannot protect one center at a time?

Each element received one of four decisions: include as-is, adapt, simplify, or set aside, with the reason recorded. Those decisions live in the Comprehensive Framework Source Register, a working document prepared with this Playbook and available to the Massachusetts League of Community Health Centers.

Consistent use of this Playbook means health centers ask the same protective questions and hold AI vendors to the same standard.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 2: Getting Started

Set Up or Strengthen Governance	Listen and Learn	Assess and Decide	Plan and Implement	Steward	Keep It Running

GOAL OF THIS SECTION

See what AI is already in your building, name who decides, and follow a 90-day path to a working governance foundation.

WHAT'S COVERED IN THIS SECTION

2.1 Current-State Assessment
Document the AI already in your building

2.2 SEEDS Decision Reasoning Maturity
See how deeply you reason through AI decisions

2.3 Minimum Viable Governance
Stand up the five components you need first

2.4 The AI Governance Hub
Form the body that decides going forward

2.5 90-Day Starter Roadmap
Follow a 90-day path to a foundation

TOOLS YOU'LL USE

Approved-Use Inventory Checklist
How do we build and maintain our AI tool inventory?

AI Policy Checklist
What does our AI policy need to cover?

AI Governance Hub Getting Started Kit
How do we set up AI governance?

Leadership Talking Points Template
What do we say about this AI decision, and who needs to hear it?

2.1 Current-State Assessment

AI is already in most health centers, and used well it can help you meet urgent challenges — if you can see it and govern it. Every health center starts here, whatever brought you to the Playbook (see the situations in Section 1.2 — How to Navigate This Playbook). Complete this foundation before entering any phase.

Work through three tools in sequence. The Approved-Use Inventory Checklist documents what AI is already in the building. The AI Policy Checklist identifies what your AI policy needs to cover. The AI Governance Hub Getting Started Kit sets up the body that governs AI decisions going forward. The picture you build here tells you where to focus in the subsections that follow. For vendor evaluation criteria, see Phase 2 — Assess and Decide (Section 4).

Start with the inventory. Document every AI tool in use, under evaluation, or recently discontinued — including tools that arrived through vendor updates and tools staff adopted on their own. For each, record approval status, who authorized it, what data it touches, and what governance is in place. Bring in the people who know: whoever manages your technology systems and contracts, the clinical or

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

operations lead who knows what departments use, and your compliance or legal support, if you have one.

Then the policy. The AI Policy Checklist walks through what your policy should address — which categories of AI tools are in scope, what standards govern their use, how new tools are surfaced and reviewed, and who provides oversight. It surfaces gaps in your current policy rather than prescribing language, so work through it with your legal and compliance support.

2.2 SEEDS Decision Reasoning Maturity

Every health center makes AI decisions. This subsection helps you see how deeply those decisions are being made today, so you know where to grow. The five maturity levels are introduced in Section 1.3 and defined in full in Appendix A: SEEDS of Innovation™; here you apply them to your own decisions.

SEEDS of Innovation™ measures reasoning maturity as five levels of depth — from Reflexive to Regenerative — assessed within each SEEDS pillar. Figure 2.1 shows the five levels at a glance.

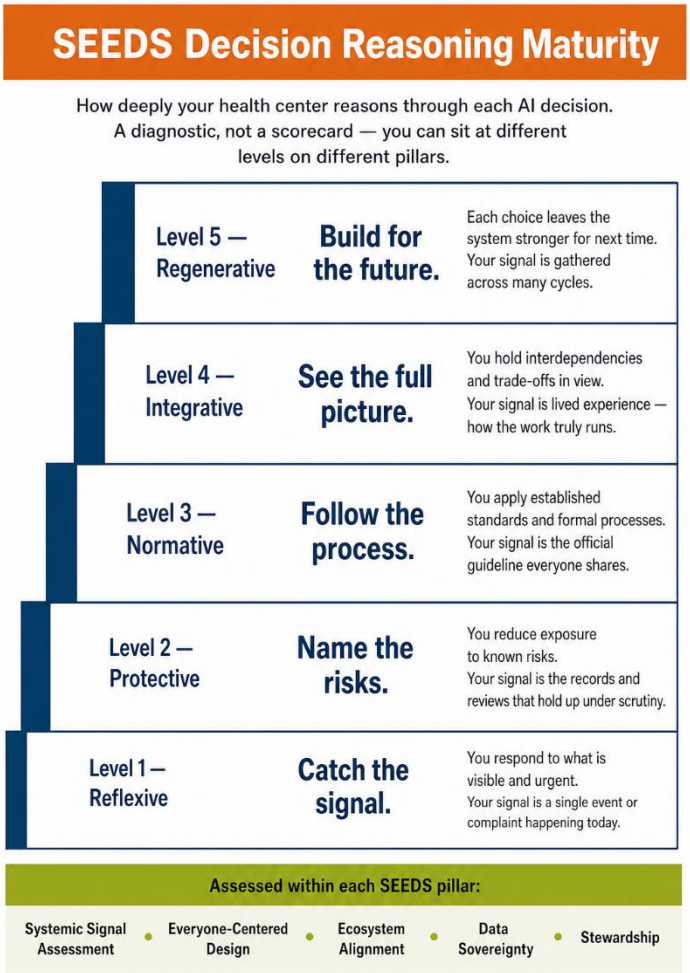

Figure 2.1. SEEDS Decision Reasoning Maturity: the five levels at a glance

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Your health center will sit at different levels on different pillars at once. That is normal, and where you land tells you where to focus first. Have your leadership team pick, for each SEEDS pillar, the level that best describes how AI decisions are actually made today, and write it down. That is your starting point.

How well your decision authority reasons matters as much as who holds it. The structure in the next subsection gives that authority a home; the depth of reasoning you just identified is what it brings to every decision.

2.3 Minimum Viable Governance

Minimum Viable Governance is the basic structure any health center can stand up, whatever its current AI maturity, staffing, or prior governance experience. It is the floor that makes the next decision governable — and what you need in place before entering any phase of the Playbook.

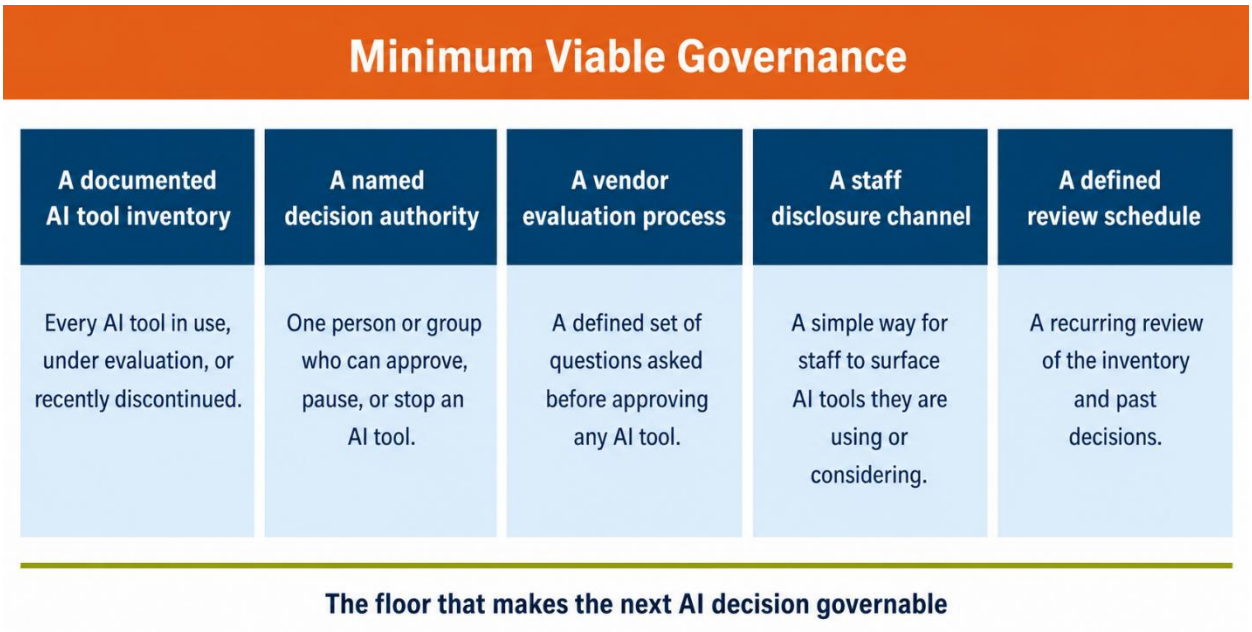

Figure 2.2. Minimum Viable Governance: the five components

Five components make up this floor. All five are essential:

- A documented AI tool inventory. Every AI tool in use, under evaluation, or recently discontinued — including vendor-update tools and tools staff adopted on their own. Start with the Approved-Use Inventory Checklist.
- A named decision authority. One person or group with clear authority to approve, pause, or stop an AI tool. The AI Governance Hub (Section 2.4) is the full version; naming one person with interim authority is the first step. The depth of reasoning this authority brings — the SEEDS maturity in Section 2.2 — is what separates a checkpoint from a decision.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- An AI policy. The written rules for AI use: what is allowed, who decides, and what every tool must meet before approval. The AI Policy Checklist in this section builds it.
- A staff disclosure channel. A simple way for staff to surface AI tools they are using or considering. Frame it as openness — staff who surface tools strengthen governance. The AI Signal Intake Form Template (Section 3) gives this channel a structured form.
- A defined review schedule. A recurring review of the inventory and past decisions. Quarterly is a reasonable start. Keep a running list of decisions still to make, each with a date to return to it.

A health center with all five can govern a new AI decision before committing to it.

2.4 The AI Governance Hub

The AI Governance Hub is a standing, cross-functional body inside your community health center (CHC). It is where AI signals — any indication that something AI-related needs attention — are reviewed, decisions are made, and outcomes are tracked. Every situation from Section 1.2 eventually routes here. It runs continuously because AI decisions are continuous, and it is cross-functional because AI touches clinical care, operations, workforce, finance, compliance, and patients. Every decision that moves through it is evaluated through the five SEEDS pillars in sequence — this is where SEEDS of Innovation™ operates inside your health center.

Who sits at the table. The Hub is built from the perspectives every AI decision needs in the room; in smaller health centers, one person may carry several. Two people keep it running day to day: the Hub Lead runs meetings and owns the agenda, and the Hub Coordinator handles logistics, scheduling, and documentation. Two functions live inside the membership rather than in separate seats — executive sponsorship (a senior leader connecting the Hub to executive decision-making) and departmental champions (connecting a specific area of the building to the Hub). Add any perspective your health center needs and does not yet see represented — for example, human resources, when an AI decision changes a role or a job description.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

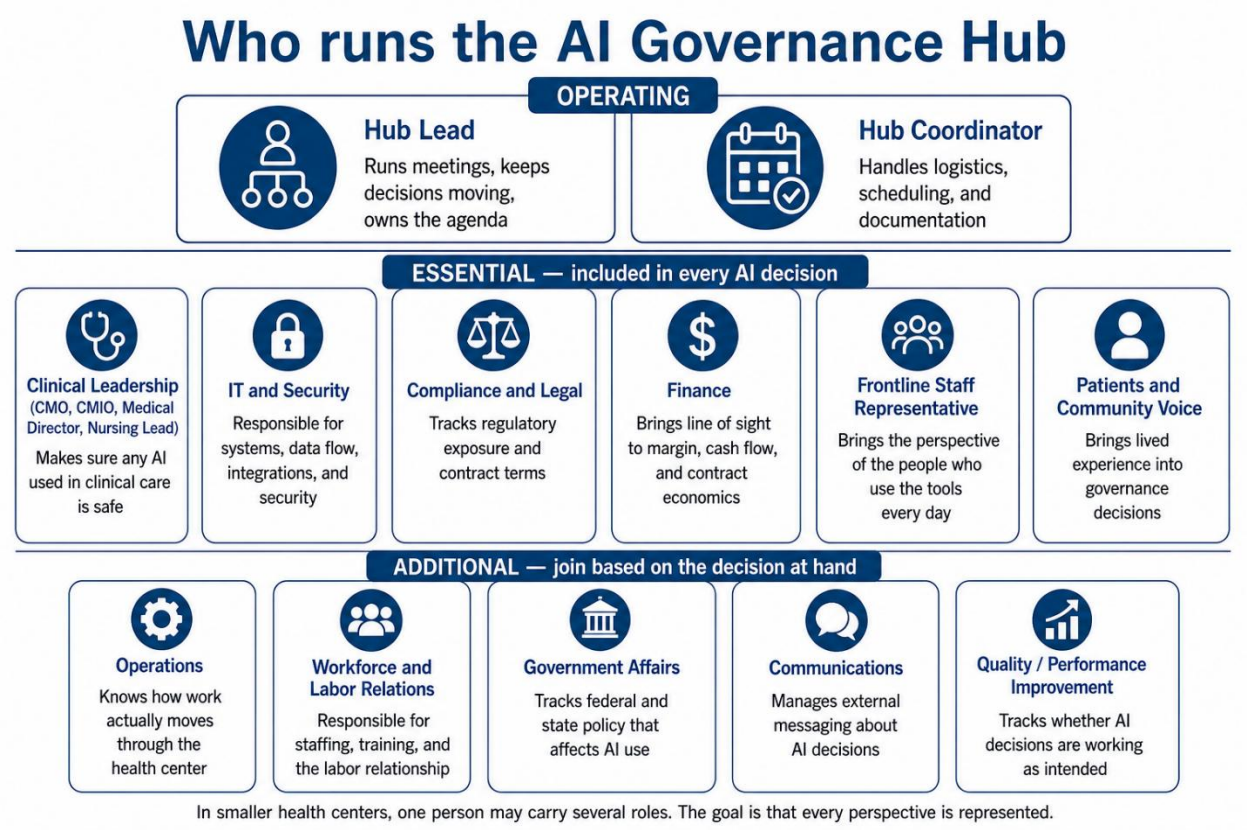

Figure 2.3. AI Governance Hub: who sits at the table and why each voice matters

How it operates. The Hub runs on two rhythms: a weekly operating meeting (30 to 45 minutes) for new signals, open decisions, and anything urgent; and a longer monthly governance meeting to examine decisions and outcomes, and work through items needing full deliberation, with a quarterly review of the full AI tool inventory. The AI Governance Hub Charter, part of the Getting Started Kit, sets minimum attendance for decisions to move forward. The Hub covers all AI-related tools, features, and processes: AI features your health record system activated through updates, staff use of AI tools whether approved or informal, vendor pitches, and patient and staff feedback.

When decisions go further. The Hub handles most AI decisions within its own authority. Legal and compliance counsel already sit in the Hub; specific matters — legally ambiguous questions, compliance events, contract review — draw on them directly. Some decisions rise to executive leadership (financial exposure beyond the Hub's authority, a significant change to a clinical service or labor relationship, high-risk assessments needing senior approval) or, through the CEO, to the board (material reputational, financial, or mission-level exposure). Whenever a decision goes further and returns, the Hub Lead records what was decided, any conditions attached, and the date, in the same place the Hub tracks every decision.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Communicating the decision. Once the Hub decides, it says so clearly. The Leadership Talking Points Template gives the Hub a consistent way to state what was decided, what it means for each audience, and who needs to hear it — staff, board, or patients.

IMPORTANT

Patient and community voice is essential for any AI decision that is patient-facing or shapes the patient experience. For tools that work behind the scenes, the Hub makes that determination deliberately. Engage patients through multiple channels: consumer board members who fill the patient seat on your Federally Qualified Health Center (FQHC) board, and direct engagement with patients facing the most barriers. Patients participating in the Hub should be compensated for their time, within applicable rules and regulations.

LEGAL NOTE

Your health center is solely responsible for the accuracy of every claim submitted to a payer (insurance company), whether or not an AI tool helped produce the documentation or the claim; that an AI tool caused an error is generally not a defense. No law currently sets a required amount of human review, but the risk of an inaccurate claim falls as human oversight increases, and the rules for AI use can vary from payer to payer. Human review of AI-assisted documentation is a governance responsibility your health center owns.

What the Mass League carries; what your health center carries. The Massachusetts League of Community Health Centers (MLCHC) works through collective action: equipping health centers with the knowledge, tools, and opportunities to adopt effective, responsible, and tailored AI — so AI helps protect access to quality care, restore time for care, and strengthen financial sustainability. Guided by its principles — Mission, People, and Community First; Collaborate and Align; Lead with Urgency and Intent — the League carries the network-level work no single health center should have to carry alone: shared vendor intelligence, shared contract and Business Associate Agreement (BAA) resources, regulatory monitoring, peer learning forums, and maintenance of this Playbook. The League convenes and equips; it does not make AI decisions for your health center.

Your health center carries the decisions: running the AI Governance Hub, your local AI tool inventory and decision authority, vendor evaluation with your own legal counsel, workforce planning and training, patient communication, and local measurement of every AI decision.

2.5 90-Day Starter Roadmap

The 90-Day Starter Roadmap gives every health center a structured path to Minimum Viable Governance, or to strengthening what is already in place. The work in Sections 2.1 and 2.2 gives you the picture; this roadmap gives you the sequence. Where you begin shapes your pace — a health center

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

with existing governance moves faster; one building from the ground up uses the full 90 days — and it works within your current staffing.

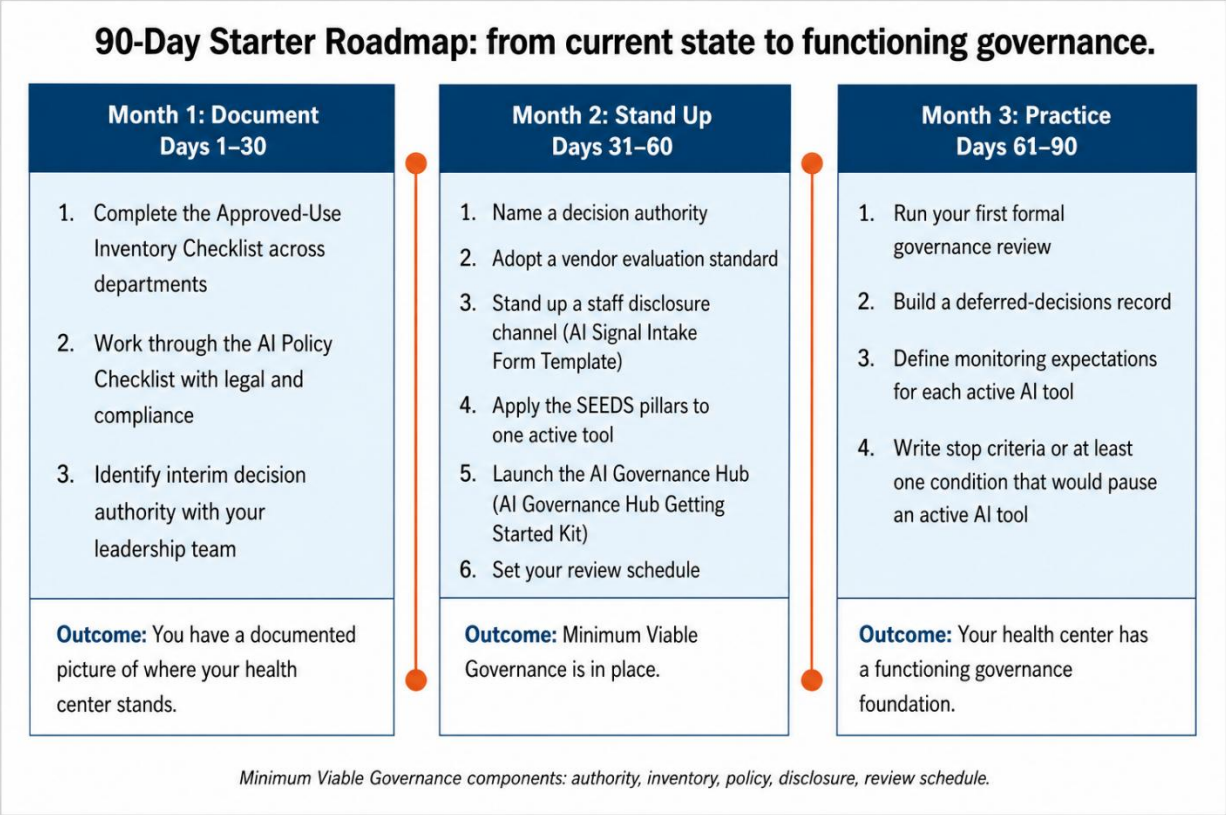

Figure 2.4. 90-Day Starter Roadmap: from current state to functioning governance

- Month 1 — Document current state. Complete the Approved-Use Inventory and see the full picture before making decisions.
- Month 2 — Formalize governance. Stand up all five components of Minimum Viable Governance: authority, inventory, policy, disclosure channel, and review schedule.
- Month 3 — Practice sustainability. Run your first formal governance review, set stop criteria for active tools, and record the decisions your team still needs to make.

At the end of 90 days, your health center has a documented current state, a functioning governance foundation, and a clear path forward.

CASE EXAMPLE: STANDING UP GOVERNANCE AT BAYSIDE COMMUNITY HEALTH CENTER

This example is hypothetical and depicts a composite community health center. The same health center appears in the case example at the end of Sections 2 through 6, so you can follow one organization through the whole process.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

The situation. Bayside Community Health Center has four sites and serves a large Portuguese-speaking patient population. During a routine software update, its health record system switched on a new AI note-taking feature. Around the same time, a few clinicians started using a free AI transcription tool they found on their own.

No one reviewed either one. Bayside has no AI governance yet. Here is how the team sets it up, one tool at a time.

Step 1 — Approved-Use Inventory Checklist *(What AI is already in our building?)*

Bayside starts by making a list. Using the checklist, the team writes down every AI tool it knows about, asks each department what they are using, and records what it finds in the inventory table: the tool and vendor, the department using it, its status, what kind of data it touches, and whether there is a signed privacy agreement (called a BAA) with the vendor.

Both tools go on the list with the status Under Review.

The checklist's last step routes each tool to a reviewer, but routing is a decision for the AI Governance Hub, and Bayside has no AI Governance Hub yet. So the team saves the routing step for later and keeps building.

Step 2 — AI Policy Checklist *(What does our AI policy need to cover?)*

Next, the team compares its current policy against the twelve elements an AI policy needs to address, answering Yes, Partially, or Not yet for each.

Three answers come back Not yet: the policy says nothing about AI built into vendor software (Element 1, Purpose and scope), no one owns keeping the approved-use list current (Element 4), and there is no rule that a person must check AI's work (Element 7, Human oversight).

For each gap, the team records an owner and a due date, just as the checklist asks. Now the policy update has a clear to-do list.

Step 3 — AI Governance Hub Getting Started Kit *(How do we set up AI governance?)*

Bayside's leadership team uses the Kit to set up its **AI Governance Hub**, the small group that will make AI decisions from now on. Part 1 is a setup checklist where every “Not yet” answer becomes the action list. Part 2 is the charter, and it walks the team through the choices one section at a time: who sits on the AI Governance Hub, how often it meets, what it can decide on its own, and how questions reach it. They name an AI Governance Hub Lead and an AI Governance Hub Coordinator and set the meeting rhythm.

Now that the AI Governance Hub exists, it completes the routing step the inventory saved for later, choosing from the checklist's three paths. The note-taking feature touches patient records, so it gets “Full review through the AI Signal Decision Assessment” (Section 4 shows how that review happens). The free transcription tool has no privacy agreement and sits outside policy, so it gets “Pause or discontinue use.”

Both decisions are written down, with the date and the conditions.

Step 4 — Leadership Talking Points Template *(What do we say about this AI decision, and who needs to hear it?)*

Pausing a tool people already use will raise questions, so the AI Governance Hub prepares its message early.

Using the template's four-step process, the AI Governance Hub drafts a short explanation, stress-tests it with the people most likely to have questions, gets executive approval, and distributes one

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

consistent message to the template's four audiences: board, staff, external partners, and patients. The message leads with why: protecting access to quality care, restoring time for care, and strengthening financial sustainability.

The outcome. In 90 days, Bayside goes from two unreviewed AI tools to a working system: a written inventory, a policy with a named owner for every gap, an AI Governance Hub that meets on a schedule, and two routing decisions on the record.

What this teaches. The four Section 2 tools build on each other in order. First you find out what you have (inventory). Then you set the rules (policy). Then you create the group that decides (the AI Governance Hub). Then you tell people what was decided (talking points). Each tool's output is the next tool's starting point. Bayside's story continues at the end of Section 3.

The full SEEDS maturity scale is in Appendix A: SEEDS of Innovation™. The public frameworks this Playbook draws on are in Appendix B: Implementation Approach Reference.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Approved-Use Inventory Checklist

This tool answers: How do we build and maintain our AI tool inventory?

Section 2 — Getting Started

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
One living record of every AI tool in your health center and its governance status. Without it, the AI Governance Hub has no reliable picture of what AI is running across the organization.	A three-step discovery process surfaces every AI tool in use and routes each one into governance. The result is a living inventory table your team updates on a defined schedule.	The AI Governance Hub coordinator maintains the master copy. Each tool owner keeps their entry current. The Hub reviews the full inventory at every quarterly meeting and acts on gaps and status changes as they arise.

Before You Begin

Gather signed vendor contracts and Business Associate Agreements (BAAs) and schedule conversations with department leads across clinical operations, billing, IT, and patient access before starting the discovery process.

Building the Inventory

Knowing what AI is running in your health center lets the AI Governance Hub assess risk, reduce duplication, and make sound decisions about adoption, training, and policy. This section guides the initial build.

IMPORTANT

Discovery is the starting condition. The inventory grows from what you find, and that is a sign the governance process is working.

Step 1: Known Tools
List the AI tools the AI Governance Hub already knows about.

Task	Yes	In Progress	Not yet

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Formally approved vendor tools from procurement records, signed contracts, or prior decisions	☐	☐	☐
Electronic health record AI features the AI Governance Hub has reviewed and documented	☐	☐	☐
Tools in active pilot with documented decisions	☐	☐	☐
Any tool with a current Business Associate Agreement (BAA) that covers AI use	☐	☐	☐

Step 2: Discovery

Surface AI tools that arrived without an AI Governance Hub decision. This is where most of the inventory work happens.

Task	Yes	In Progress	Not yet
Survey department leads (clinical, operations, billing, IT, patient access) about what AI tools or AI-enabled features their teams use	☐	☐	☐
Review vendor release notes and feature announcements for the past 12 months for AI capabilities that may have activated automatically	☐	☐	☐
Audit vendor contracts and Business Associate Agreements (BAA) for AI-specific terms	☐	☐	☐
Ask frontline staff about consumer AI tools they use for work tasks (drafting communications, summarizing notes, translating content)	☐	☐	☐
Confirm whether staff have received training and updated standard operating procedures or workflows for the AI tools their teams use	☐	☐	☐

Step 3: Routing

Once the AI Governance Hub is established, route each newly surfaced tool through the AI Signal Intake Form Template in Phase 1 — Listen and Learn (Section 3). Until then, log what you find in the inventory table below and return to routing after Section 2 setup is complete. The AI Governance Hub routes each tool to one of three paths:

- ☐ Full review through the AI Signal Decision Assessment when the tool use case, data access, or population impact warrants deeper assessment.
- ☐ Expedited approval when the tool clearly meets existing criteria and the AI Governance Hub can document the basis.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- ☐ Pause or discontinue use when the tool falls outside the health center AI policy and approved-use framework.

Once routing is underway, build the inventory table below with all known and newly surfaced tools.

The Inventory Table

Maintain the inventory in a spreadsheet (such as Excel or Google Sheets) so your team can sort, filter, and update it as a living record. Because entries can reference protected health information, data types, and BAA status, store the file where your health center controls access.

The example rows show how to fill in each column. Replace them with your health center's actual tools, and add rows as needed. Health centers may add columns for context-specific tracking (renewal dates, account contacts, compliance notes) while keeping the eight core columns intact.

Tool / Vendor	Department	Status	Authorized By	Date	Data Type	BAA Status	Next Review
AI Ambient Documentation / [Vendor]	Primary Care	Approved	Hub Lead, [Name, Title]	03/15/2026	PHI: real-time encounter audio and transcript	In Place	09/15/2026
Language Access AI / [Vendor]	Patient Access, Behavioral Health	In Pilot	Hub Lead, [Name, Title]	01/10/2026	Patient-facing content, de-identified notes	In Place	07/10/2026

Maintaining and Using the Inventory

The inventory is only as useful as it is accurate. Update it when any of the following occurs:

- ☐ At every quarterly AI Governance Hub review: confirm all entries are accurate
- ☐ When a new AI tool is approved: add a row to the inventory
- ☐ When a tool status changes (approved, in pilot, under review, decommissioned): update the Status column
- ☐ When a BAA is signed, updated, or expires: update the BAA Status column
- ☐ When the AI Governance Hub completes a stewardship review: update the Next Review date

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

☐ When a vendor announces a material change (model update, ownership change, data-handling change): set status to Under Review

The inventory also surfaces conditions that require the AI Governance Hub to act:

- A status change to Under Review or Decommissioned triggers AI Governance Hub reassessment or the decommissioning process in the Stewardship Review Checklist, Phase 4 — Steward (Section 6).
- Tools in active use absent from the inventory surface a gap for the AI Governance Hub to address.
- Patterns across multiple tools (accessing protected health information without clear data residency rules, missing BAAs, missing designated owners) inform AI policy updates and training priorities.

KEY TAKEAWAYS BY ROLE

Leadership

- The inventory gives you a complete, verified picture of what AI tools are in use across the organization before you make resource, budget, or risk decisions.
- Where the inventory shows gaps (tools in use with missing BAAs, unassigned owners, or unresolved data-handling questions), those gaps are the AI Governance Hub's first action items.

AI Governance Hub members

- You own the inventory. Assign row owners, drive discovery conversations with department leads, and schedule the quarterly review cycle before you start.
- Treat the inventory as a living record. A complete and current inventory is the foundation for every policy, training, and vendor decision the Hub makes.
- Route every newly surfaced tool through the AI Signal Intake Form Template (Section 3) once the AI Governance Hub is established.

IT & Security

- Your role in discovery is active. Department leads depend on you to surface tools embedded in existing systems, vendor integrations, and infrastructure that sit outside departmental visibility.
- Flag every tool that touches protected health information for BAA review, even when the tool appears low-risk. That determination belongs to the AI Governance Hub.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

Before evaluating whether AI is the right solution, you need a clear picture of what is already in use. The inventory gives you that baseline. After completing or reviewing the inventory, move to Phase 1 — Listen and Learn (Section 3) to assess whether AI is the right path for the new need.

→ Evaluating a specific AI tool?

Log a newly approved tool before it goes into use. Complete the governance assessment in Phase 1 — Listen and Learn (Section 3) first, then return here to add the tool to the inventory once a decision is reached.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

→ Formalizing an existing AI tool?

This is the inventory's most common entry point. Log the tool, assign a provisional owner, and route it through the AI Signal Intake Form Template (Section 3) for assessment. The discovery process in Building the Inventory is designed for exactly this situation.

→ Reassessing a changing AI tool?

Update the tool status to Under Review and route it back to Phase 1 — Listen and Learn (Section 3) for reassessment. The original inventory row stays in place, and the status column carries the change.

→ Checking an approved AI tool?

Confirm the tool's designated owner, BAA status, data access scope, and Next Review date are current before proceeding. Then move to the Stewardship Review Checklist in Phase 4 — Steward (Section 6).

→ Replacing or exiting an AI tool?

Update the status to Decommissioned, document the exit rationale in the Notes column, and flag data migration or contract exit obligations for the AI Governance Hub. To replace the tool, run the Evaluating a specific AI tool path for the incoming tool and return here to add the new entry after approval.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI Policy Checklist

This tool answers: What does our AI policy need to cover?

Section 2 — Getting Started

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A twelve-element checklist covering what each community health center (CHC) AI policy needs to address. Use it to build a new policy or evaluate an existing one.	Work through each element in order. Confirm whether the current policy covers it, select Yes, Partially, or Not yet, and record an owner, due date, and note for any gaps. Each element points to the Playbook tools and sections that put it into practice.	The AI Governance Hub (the cross-functional team that oversees AI decisions) completes this checklist as part of standing up minimum viable governance. Legal counsel and the compliance lead review the resulting policy. The AI Governance Hub Lead signs off before adoption.

Before You Begin

Gather a current AI policy draft, or confirm the intent to write one. A list of AI tools currently in use is also useful; the Approved-Use Inventory Checklist supports this step. Have current vendor contracts and Business Associate Agreements (BAAs) on hand for Elements 5 and 8.

Select Yes when the policy fully covers the element as described. Select Partially when coverage exists but is incomplete. Select Not yet when the element is absent from the current policy. Record an owner, due date, and short note for every Partially or Not yet item.

IMPORTANT

This checklist guides what a policy can cover. Each health center decides which elements it adopts and how, based on its own legal and operational context. Use it to confirm coverage and surface gaps, then set the owner, due date, and note for each gap.

LEGAL NOTE

This checklist is a general governance resource, not legal advice. Each health center is solely responsible, in conjunction with its own independent legal counsel, for ensuring its AI uses and policies are permissible under applicable law. AI is a fast-developing technology, and regulatory

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> and policy changes are expected. CHCs should confer with legal counsel to determine how recent or anticipated changes affect their policy. The specific language, scope, and applicability of each element are the health center's to determine.

The twelve elements

POLICY ELEMENT	YES	PARTIALLY	NOT YET
1. Purpose and scope. The policy names the categories of AI use it governs: clinical tools, administrative tools, patient-facing tools, and AI features built into existing vendor software. It states what is in scope, what is out of scope, and how the scope is updated when new types of tools are introduced. Each person who uses or approves AI tools on behalf of the health center falls within the policy's scope, including contracted staff, vendors, consultants, and board members in their governance role. Staff use only approved AI tools for health center work, and keep patient health information (PHI) and other personal information within those approved tools. *Playbook connection: applies from the first AI signal through each phase of the SEEDS of Innovation™ decision process.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
2. Decision authority. The policy names who holds authority to approve AI tools for use. The AI Governance Hub typically holds core approval authority, with specific roles named for clinical, contract, and data decisions. It defines the escalation path for decisions that exceed the AI Governance Hub's authority. Examples include tools with significant patient safety implications, tools that affect contractual obligations, and tools that require board notification. *Playbook connection: Section 2 — Getting Started; Section 2.4 — The AI Governance Hub.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
3. Approval criteria. The policy states the minimum criteria any AI tool must meet before it is approved for use: clinical and operational evidence, data governance requirements, vendor accountability, staff readiness, and a documented human oversight plan. It references the AI Signal Decision Assessment as the	☐	☐	☐

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

POLICY ELEMENT	YES	PARTIALLY	NOT YET
standard process for applying these criteria to each tool under consideration. *Playbook connection: Phase 2 — Assess and Decide (Section 4), AI Signal Decision Assessment.* If Partially or Not yet: Owner _____ Due _____ Note _____			
4. Approved-use list maintenance. The policy requires that every AI tool in use at the health center appears on the Approved-Use List. The list records each approved tool's scope, data permissions, conditions of use, assigned owner, and next review date. The policy names who maintains the list, how often it is reviewed, and what triggers a status change. *Playbook connection: Section 2 — Getting Started, Approved-Use Inventory Checklist.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
5. Patient information and data handling. The policy defines how patient health information (PHI), meaning information that identifies patients and relates to their health or care, is handled in AI tools. It requires a Business Associate Agreement (BAA) with any vendor who handles PHI, names data retention and deletion requirements, and identifies applicable federal and state privacy laws. Vendor data use rights and model training rights are addressed directly: the policy states what the health center permits and what it prohibits with respect to its patient data. *Playbook connection: Phase 2 — Assess and Decide (Section 4.1, Vendor Evaluation); Phase 4 — Steward (Section 6).* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
6. Patient notice and consent. The policy defines when patients are informed that AI is being used in their care and how that notice is delivered in plain language. Language access requirements are named: patients who speak a language other than English receive notice in a language they understand. For tools that record patient conversations, generate patient-facing content, or make recommendations that affect patient care, the policy states explicit consent requirements. *Playbook connection: Phase 1 — Listen and Learn (Section 3.3, Patient, Staff, and Board Voice).*	☐	☐	☐

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

POLICY ELEMENT	YES	PARTIALLY	NOT YET
If Partially or Not yet: Owner _____ Due _____ Note _____			
7. Human oversight requirements. The policy names the conditions under which a qualified person reviews, approves, or overrides AI output before it is used. This includes clinical documentation, recommendations that affect patient care, and decisions that affect patient benefits or services. It names the right of any staff member to override AI output without penalty, meaning every performance measure and operational process allows adequate review of AI output before it is acted on. When patient information reaches an unapproved tool, or AI output is used without authorization, staff report it through the incident process in Element 11, and consequences for policy violations follow the health center's existing personnel and compliance policies. *Playbook connection: applies across all phases; required human review for clinical use is defined in Phase 2 — Assess and Decide (Section 4).* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
8. Vendor and procurement requirements. The policy names what the health center requires from the companies that provide AI tools before a contract is signed. Required documentation includes a Business Associate Agreement (BAA), evidence of regulatory clearance, documentation of bias testing for the patient population the health center serves, and evidence that the tool produces the claimed results for communities with barriers to care. Contract terms cover data ownership, data return or deletion when the contract ends, and the vendor's obligation to notify the health center of material changes to the tool, including model updates, ownership changes, and changes to data handling. *Playbook connection: Phase 2 — Assess and Decide (Section 4.1, Vendor Evaluation). The Vendor Evaluation Checklist is the operational tool.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
9. Staff training and competence. The policy names what training staff complete before using any approved AI tool and how that training stays current. Training requirements are role-specific: clinical staff, administrative staff, and AI Governance Hub members	☐	☐	☐

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

POLICY ELEMENT	YES	PARTIALLY	NOT YET
each have different responsibilities and need different preparation. Training is offered in the languages the workforce uses. Staff complete required training before any tool is put in place and receive updated training when a tool changes materially. *Playbook connection: Phase 3 — Plan and Implement (Section 5.3, Workforce Lens and Role Changes). The AI Implementation Readiness Checklist is the operational tool.* If Partially or Not yet: Owner _____ Due _____ Note _____			
10. Monitoring and stewardship. The policy names the AI Governance Hub's responsibility to monitor AI tools after they are put in place. It sets the default review cadence, at minimum quarterly, and names the conditions that trigger an unscheduled review, including patient safety concerns, significant performance changes, and material changes by vendors. *Playbook connection: Phase 4 — Steward (Section 6), Stewardship Review Checklist.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
11. Incident reporting and response. The policy defines how AI-related incidents are surfaced, reported, and addressed. This includes patient safety concerns, data events, and material changes made by vendors. A reporting-safe standard is named: a staff member who reports a concern in good faith is protected from discipline for the report itself. Response timelines, notification requirements, and escalation paths are defined by the severity of the incident. *Playbook connection: Phase 1 — Listen and Learn (Section 3), AI Signal Intake Form Template.* If Partially or Not yet: Owner _____ Due _____ Note _____	☐	☐	☐
12. Policy review and maintenance. The policy defines how it is kept current. At minimum, the AI Governance Hub reviews the policy annually. Review is also triggered when significant regulatory changes, vendor contract changes, or operational changes affect how AI is used. The AI Governance Hub Lead is accountable for ensuring the review happens, for documenting what changed and why, and for obtaining new executive approval	☐	☐	☐

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

POLICY ELEMENT	YES	PARTIALLY	NOT YET
after any material revision. The version number and effective date are updated after each review. *Playbook connection: Section 7 — Keeping the System Running.* If Partially or Not yet: Owner _____ Due _____ Note _____			

KEY TAKEAWAYS BY ROLE

Leadership

- The AI policy is the executive commitment that governs each AI decision at the health center. Approve the final policy and ensure it carries board-level visibility.
- Require that each AI tool in use falls within this policy's scope before it operates.

AI Governance Hub members

- Own this checklist. Work through each element, confirm coverage, and assign owners and timelines for any gaps.
- Return to it at each annual policy review and whenever a material change triggers a review.

IT & Security

- Confirm that Elements 5 (data handling), 8 (vendor requirements), and 11 (incident reporting) align with existing security policies and BAA language.
- Flag any element where current technical infrastructure creates a gap the policy must acknowledge.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

Before exploring any tool, confirm your health center has a policy that covers the twelve elements. This checklist is part of minimum viable governance in Section 2 — Getting Started.

→ Evaluating a specific AI tool?

Confirm your policy has the decision authority, approval criteria, and vendor requirements in place first. Elements 2, 3, and 8 are most relevant. Then move to Phase 2 — Assess and Decide (Section 4).

→ Formalizing an existing AI tool?

The policy is your first reference point. Use this checklist to confirm whether what is already in use falls within the policy's scope and whether the policy has the elements to govern it. Then move to Phase 1 — Listen and Learn (Section 3) to route the signal.

→ Reassessing a changing AI tool?

When a vendor adds features, changes data handling, or expands scope beyond the original decision, check whether the current policy still covers the new conditions. Where a gap appears, record it in the Owner, Due, and Note fields.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

→ Checking an approved AI tool?

Before running the check, confirm the policy's monitoring cadence (Element 10) and review triggers (Element 12) are in place. Then move to Phase 4 — Steward (Section 6).

→ Replacing or exiting an AI tool?

The policy governs what happens when a tool exits: data return or deletion (Element 5), contract wind-down (Element 8), and approved-use list updates (Element 4). Confirm these are in place, then move to Phase 4 — Steward (Section 6) for exit planning.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI Governance Hub Getting Started Kit

This tool answers: How do we set up AI governance?

Section 2 — Getting Started

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A setup checklist and charter template for standing up your AI Governance Hub, the body behind every AI decision.	Work the setup checklist first, then complete the charter: who serves, what the Hub decides, and how it operates.	This Kit builds the Hub itself, naming members, setting the meeting rhythm, decision rules, quorum, and escalation.

Before You Begin

Your leadership team completes this Kit together. Before you start, name an executive sponsor and confirm that standing up the AI Governance Hub is your organization's next step.

PART 1 — SETUP CHECKLIST

One continuous checklist. Work top to bottom; each item is a confirmation your leadership team makes before the Hub convenes. A "Not yet" answer is your action list.

Setup item	Yes	Not yet
Sponsorship. A named executive sponsor is identified and has agreed to the role. STOPGATE	☐	☐
The executive sponsor accepts accountability for Hub decisions that reach their level.	☐	☐
The executive sponsor is ready to champion the Hub's work to the board and senior leadership team.	☐	☐
Meeting rhythm. A rhythm is set: a weekly operating meeting (30 to 45 minutes) for signals, open decisions, and urgent items, and a monthly governance meeting for outcomes and policy, with a quarterly review of the full AI tool inventory. Smaller health centers may combine the two. Record your schedule in charter Section D.	☐	☐

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Setup item	Yes	Not yet
A recurring calendar hold is scheduled for all Hub members.	☐	☐
A standing agenda format is set (signals review, active tool check, decisions required, next steps).	☐	☐
A notetaking and documentation process is identified.	☐	☐
Perspectives. The perspectives in the membership table (charter Section C) are identified and confirmed, including the Hub Lead and Hub Coordinator.	☐	☐
Decision authority. Decision-authority thresholds are agreed and recorded in charter Section E. STOPGATE	☐	☐
Quorum and voting. Quorum and voting rules are agreed and recorded in charter Section F.	☐	☐
Charter. The charter (Part 2) is completed with current member names, roles, and decision rules.	☐	☐
The completed charter is shared with all Hub members.	☐	☐
The completed charter is shared with executive leadership.	☐	☐
A process for updating the charter when membership or rules change is identified. STOPGATE	☐	☐
The AI Governance Hub is ready to convene. STOPGATE	☐	☐

PART 2 — HUB CHARTER TEMPLATE

Fill each field. Italic text shows what to enter; Example shows a model answer you can adapt.

Section A — Health Center Information

Health center name	
Charter date	
Charter version	1.0
Last reviewed	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section B — Purpose and Scope

Purpose statement	*Adapt this statement.* **Example:** The AI Governance Hub is [Health Center Name]'s standing decision-making body for AI decisions. It reviews AI signals, evaluates new and active tools, and confirms each AI decision reflects our values, protects our patients, and fits our operations.
Scope of authority	*List what the Hub governs.* **Example:** All AI tools used by staff; all AI used in or connected to patient care; all AI-related vendor relationships.
Outside Hub review	*Name routine features that fall outside review until they add AI functions affecting clinical decisions or patient data.* **Example:** standard calendar scheduling; document formatting.

Section C — Hub Membership

Role	Perspective they bring	Name	Confirmed?	Notes
Executive sponsor	Organizational authority; final escalation point; board and leadership liaison		☐	
Hub Lead (assigned from existing members)	Runs meetings, owns the agenda, manages the signal log		☐	
Hub Coordinator (assigned from existing members)	Scheduling, documentation, follow-up		☐	
Clinical leadership (CMO, Medical Director, or equivalent)	Care delivery, clinical workflow, patient safety		☐	
Nursing leadership	Nursing practice, care-team workflow, frontline clinical impact		☐	
Operations leadership (COO or equivalent)	Staff capacity, day-to-day impact		☐	
IT or Health Informatics	Technical infrastructure, data systems, integration		☐	
Compliance or Legal	Regulations and policy, liability, HIPAA and BAA requirements		☐	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Role	Perspective they bring	Name	Confirmed?	Notes
Finance	Budget, vendor contracts, financial risk		☐	
Human Resources, Workforce, and Training	Staffing, role changes, reskilling and training plans		☐	
Frontline staff representative (nurse, MA, patient navigator, or care coordinator)	How AI tools land in practice; staff impact		☐	
Patient and community member	Lived experience of care; access and community benefit		☐	
Quality / Patient Safety	Outcome data, safety monitoring, improvement cycles		☐	
Government Affairs / Policy	Regulatory and payer landscape, state policy, advocacy		☐	
Communications	Patient and staff messaging, AI disclosure, public trust		☐	
Other: _____	A perspective specific to your health center		☐	

In smaller health centers one person may carry several perspectives; the goal is that every perspective is represented. HIPAA is the Health Insurance Portability and Accountability Act; BAA is a Business Associate Agreement.

Section D — Meeting Structure

Meeting rhythm	*Record your cadence.* **Example:** weekly operating meeting (30 to 45 minutes); monthly governance meeting (longer).
Standard meeting length	
Format (in-person, virtual, or hybrid)	
Standing agenda format	*Set your standing items.* **Example:** signals review; active tool check; decisions required; next steps.
Notetaking responsibility	
Where notes are stored	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section E — Decision Authority

Record your authority and escalation for each decision type.

Decision type	Your entry (Example shown)
Approve a new AI tool for limited use	**Example:** Full Hub vote; executive sponsor if board approval is required.
Approve a new AI tool for organization-wide use	**Example:** Hub recommends; executive sponsor confirms; board if capital or policy impact.
Pause or remove an active AI tool	**Example:** Hub Lead and executive sponsor decide together.
Approve a vendor contract or regulatory agreement	**Example:** Hub reviews; Finance and Legal confirm.
Respond to a regulatory inquiry or audit related to AI	**Example:** Hub Lead coordinates; Legal and executive sponsor as required.

Section F — Quorum and Voting Rules

Record your rule for each.

Rule	Your entry (Example shown)
Quorum	**Example:** a majority of Hub members present, including at least one of executive sponsor, clinical lead, or operations lead.
Quorum for vendor or regulatory review	**Example:** the Compliance or Legal perspective is present.
Tie or split vote	**Example:** Hub Lead and executive sponsor decide together.
Urgent decision between meetings	**Example:** Hub Lead consults the executive sponsor; the decision is documented and confirmed at the next meeting.

Section G — Signal Routing

> **IN PLAIN TERMS**
>
> Signal: any indication, from a person, a system, a vendor, a regulator, or a peer organization, that something AI-related needs attention.

Confirm or adjust how each signal type routes.

Signal type	Your entry (Example shown)
New AI tool under consideration	**Example:** Phase 2 — Assess and Decide (Section 4).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Signal type	Your entry (Example shown)
AI tool already in use without review	**Example:** Phase 2 — Assess and Decide (Section 4), retroactive review.
Active tool performance concern	**Example:** Phase 4 — Steward (Section 6).
Vendor-initiated change to an active tool	**Example:** Phase 3 — Plan and Implement (Section 5) or Phase 4 — Steward (Section 6), depending on scope.
Regulatory or compliance signal	**Example:** Hub Lead notifies Compliance or Legal right away; the Hub reviews at the next meeting.

Section H — Charter Review and Updates

Annual review date	
Who initiates the review	
Who approves updates	
Where the current charter is stored	
How members are notified of updates	

KEY TAKEAWAYS BY ROLE

Leadership

- Name an executive sponsor before anything else; the Hub operates on authority anchored at the top.
- Review the completed charter before the first meeting; your agreement on decision authority and quorum rules is what makes the Hub work.
- Use the AI Signal Decision Assessment (Phase 2 — Assess and Decide, Section 4) for decisions that call for a full SEEDS-pillar review.

AI Governance Hub members

- Work the setup checklist in order; each item builds the foundation for the next.
- Complete the charter before the Hub convenes; it is the reference point for every future decision.
- Update the charter whenever membership or rules change.

IT & Security

- Confirm your seat before the first meeting; AI decisions affecting infrastructure and patient data call for your perspective.
- Flag any vendor-initiated tool change as a signal as soon as you see it; it routes to Phase 2 — Assess and Decide (Section 4) for reassessment, or to Phase 4 — Steward (Section 6) when the tool is already live and under regular review.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

Stand up your Hub with this Kit first, then use Phase 1 — Listen and Learn (Section 3) to capture and assess signals.

→ Evaluating a specific AI tool?

Your Hub needs to be in place before evaluation begins. Complete the setup checklist, then move to Phase 2 — Assess and Decide (Section 4) and the AI Signal Decision Assessment.

→ Formalizing an existing AI tool?

Stand up your Hub with this Kit, then use Phase 2 — Assess and Decide (Section 4) for a review of tools already in use.

→ Reassessing a changing AI tool?

Your AI Governance Hub reviews the change using Phase 2 — Assess and Decide (Section 4). Update charter Section G if signal routing or decision authority shifts.

→ Checking an approved AI tool?

Your Hub runs the review using Phase 4 — Steward (Section 6). Record findings in the meeting record named in charter Section D.

→ Replacing or exiting an AI tool?

Your Hub makes this decision using Phase 4 — Steward (Section 6). The thresholds in charter Section E set whether escalation to the executive sponsor or board is required.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Leadership Talking Points Template

This tool answers: What do we say about this AI decision, and who needs to hear it?

Section 2 — Getting Started

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
Leadership messages for two jobs: introducing AI and your AI governance approach to the organization, and communicating specific AI decisions. Four audiences: board, staff, patients, and external partners.	Three parts: a four-step process from draft to distribution; four audience templates with worked examples built in; and common failure modes to avoid.	The AI Governance Hub drafts, stress-tests, and revises messages when governance is introduced or an AI decision changes. Executive leadership approves before distribution.

Before You Begin

What you need before you start: a clear description of what you are communicating, whether the launch of your AI governance approach or a specific AI decision, including what the tool does, who it affects, and what is changing; the audiences who need to hear it, since some messages reach only one or two of the board, staff, external partners, and patients; and your supporting documentation, including the SEEDS-based decision record, for reference as you draft.

Fill in every bracketed placeholder before distributing to any audience.

Part 1: The Process

Effective talking points are short, clear, and approved by appropriate executive leadership before use. This four-step process gives the AI Governance Hub a clear path from first draft to distribution.

Step 1: Develop

The AI Governance Hub develops a draft whenever any of the following happens:

- Your health center is introducing its AI governance approach and standing up the AI Governance Hub.
- A new AI tool is approved, paused, or stopped.
- A peer community health center (CHC) reports a notable AI experience, such as a vendor failure, a regulatory action, a measured benefit, or a measured harm.
- A regulatory or payer (insurance company) change affects how AI can be used.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- Staff or board members raise a recurring question that still needs a talking point.

Each draft starts with two questions for every audience: What does this audience need to know? And what does this audience need from us? The answers drive the content.

Step 2: Stress-Test

Before talking points go live, the AI Governance Hub runs a structured review with the people most likely to have questions. Include at least one reviewer from each group:

- A frontline staff member affected by the tool.
- A patient-facing role, such as a community health worker, patient advocate, or community board member.
- A clinical or operational lead.
- A compliance or legal lead.
- An information technology or security lead when the message makes claims about data handling, privacy, or security.

Ask each reviewer the same five questions:

- What would a thoughtful skeptic say is missing?
- Does any part of this sound defensive or dismissive?
- Can we prove or measure everything we are claiming here?
- If a patient or frontline staff member reads this, what might land differently than intended?
- If this were quoted at a board meeting in six months, would we still stand behind it?

When a situation requires immediate communication, run an abbreviated review with at least a clinical or operational lead and a compliance lead, and document what was tested.

Step 3: Approve

After the stress-test, the AI Governance Hub revises and appropriate executive leadership approves. These talking points represent where your health center stands on AI, so approval sits at the executive level and aligns the message with your organization's AI strategy.

Step 4: Distribute and Use

Approval covers the health center's position and the structure of the message. Local teams fill the bracketed placeholders with their own details, and the approved position stays consistent across all four audiences. Flag any situation that still needs talking points and bring it back to the AI Governance Hub.

Part 2: Four Audiences

Each audience gets a one-page template with a worked example built in. The italic text is the instruction; the regular text beginning Example is a worked example you can adapt. Customize the bracketed placeholders for the specific AI decision and keep the structure intact.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Lead with the why. The Massachusetts League of Community Health Centers (MLCHC) AI strategy frames the purpose every message connects to: protect access to quality care, restore time for care, and strengthen financial sustainability. Three principles guide it: Mission, People, and Community First; Collaborate and Align; and Lead with Urgency and Intent. MLCHC executive leadership has called on health centers to use their collective leadership and buying power so AI is designed for the patients and staff health centers serve. Connect every message to the problem being solved, and lead with that problem rather than the technology.

Adjust your tone for each audience. The structure of what to say stays the same across all four audiences; the register changes.	
Board	Lead with patient outcomes and stewardship of resources. Use concrete numbers where possible. Keep technology in the background.
Staff	Be specific, name what is still being learned, and follow through on every commitment. Frame AI as a decision being made with their input.
External partners	Be detailed and operational for payers. Lead with mission for community partners. Share results honestly with peer health centers, including what you learned and would do differently.
Patients	Use the simplest language available. Communicate before the tool is in use. For tools that touch clinical visits, make sure patients have a real choice. Honor language access needs. Keep the path to questions short and clear.
To collect patient feedback on AI use after communicating, see the Patient AI Feedback Form Template in Phase 1 — Listen and Learn (Section 3).	

AUDIENCE 1: BOARD Use this when briefing your health center's board. Most board items get five minutes or less, so this is built for that.	
[Decision or topic]	*One sentence stating the AI decision, signal, or topic the board needs to know about.* **Example:** We are beginning a limited test of an ambient AI scribe, a tool that listens during a visit and drafts the clinical note for the provider to review, with eight providers in primary care and behavioral health starting next month.
[What we are doing and why]	*Two to three sentences. Lead with the patient or operational outcome and the problem being solved, and connect it to your strategy.* **Example:** Providers tell us documentation is the largest source of after-hours work. This test measures whether the tool restores time for care while holding note quality, in line with our strategy to restore time for care.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

[How this fits our governance]	*One sentence naming the SEEDS pillar that surfaced the issue and the role of the AI Governance Hub.* **Example:** Our AI Governance Hub reviewed the tool against the five SEEDS pillars, the decision method we use so AI choices stay rooted in mission, people, and community, and approved it as a limited test with defined stop criteria.
[Risks we are watching]	*Name the top one or two risks. Boards trust leaders who name risks directly.* **Example:** We are watching whether the tool captures the social and behavioral context our patients bring, and whether providers keep reviewing every note as the tool becomes routine. Both are measured at 30, 60, and 90 days.
[What we will report back]	*One sentence on what the board will hear next and on what cadence.* **Example:** The board will receive a 90-day review with a continue, adjust, or stop recommendation.
[Environmental and resource considerations]	*Optional. One sentence on whether the vendor provided information on energy use or data center location during your evaluation.* **Example:** We asked the vendor about the tool's energy use and data center location, and will share what we learn.

AUDIENCE 2: STAFF Staff communication on AI is already happening. This gives leadership a structure to lead it.	
[What is changing]	*Plain language: what AI tool, in what part of the work, starting when.* **Example:** Starting next month, eight providers begin a limited test of an ambient AI scribe that drafts a clinical note for the provider to review and edit before signing.
[What this means for your work]	*Be specific about responsibility, and frame staff as in control.* **Example:** If you are one of the eight providers, you remain in control of and responsible for every note. Read each one and confirm it is accurate and complete before it becomes part of the chart. If you are outside the test, your work stays the same.
[What stays the same]	*Name it. This matters as much as what is changing.* **Example:** Patient relationships, clinical judgment, and your role stay the same. The tool supports documentation only.
[How we are protecting and investing in the workforce]	*Name the employer commitment, the AI literacy education plan, and any reskilling plan and budget. Keep only the promises you can keep.* **Example:** As a major employer in the communities we serve, we are investing in AI literacy education for all staff and pairing any role changes with a named reskilling plan. Any time the tool saves stays with the provider.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

[Where to bring concerns]	*Name a specific path, including for values-based concerns.* **Example:** Bring questions about process, accuracy, or patient experience to your clinical lead. Bring concerns about whether we should use AI in clinical care to the chief medical officer or chief executive officer directly. Both are tracked.

IMPORTANT

Some staff may raise concerns about AI on personal or professional values grounds, such as its environmental cost, its effect on how people think, its use in other industries, or its fit with their role. These concerns are legitimate and deserve a direct response. Staff need a named channel to raise them, a leader who will listen and respond, and clear language about which decisions are open to input and which have been made. Section 5 (Phase 3 — Plan and Implement) develops the full framework for these conversations.

AUDIENCE 3: EXTERNAL PARTNERS External partners include community partners, payers, advocacy organizations, peer health centers, and MLCHC. Each brings a different question, and the structure stays the same.	
[Our current AI use]	*A short summary, drawn from your Approved-Use Inventory Checklist.* **Example:** Our AI tools in active use are listed on our Approved-Use Inventory Checklist and reviewed at regular intervals.
[Our governance approach]	*Two to three sentences naming the AI Governance Hub and how you evaluate AI tools, including how patient and staff needs weigh in the selection.* **Example:** Our AI Governance Hub reviews every tool against the five SEEDS pillars before approval, weighing patient and staff needs in the selection itself.
[What this means for our work together]	*Specific to the partner, including how patient and staff needs were weighed in selecting the tool.* **Example:** For payers: how AI affects coding, documentation, and claims. For community partners: how it affects patient experience and language access. For peer health centers: what we are learning. For MLCHC: what signals to feed back to the network.
[Where to bring questions or concerns]	*Name a contact path. External partners need to know who to talk to.* **Example:** Questions about our use of AI can go to [named path].
[Our environmental governance approach]	*Optional, for community partners and advocacy organizations: whether your health center weighs environmental impact in AI vendor evaluation, consistent with your community mission.* **Example:** As part of our vendor evaluation, we ask AI vendors about the environmental footprint of their tools, including data center location and energy commitments.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> **LEGAL NOTE**
> Any payer-facing statements about how your health center handles AI in claims, documentation, or coding must be reviewed by your legal and compliance team before use. This template is a communication guide only.

AUDIENCE 4: PATIENTS Patients have a right to know when AI is involved in their care. Communicate early, in plain language, with a clear path for questions. (Send in the patient's preferred language on file.)	
[What is changing in your care]	*Plain language: what the AI tool does, during which part of the visit, starting when. One to two sentences.* **Example:** Starting [date], some providers at our health center will use an AI tool during visits. The tool listens during your appointment and helps write the visit notes.
[What this means for you]	*Name who reviews what the tool produces before it affects care.* **Example:** Your provider reads everything the tool writes before it goes into your health record, and your provider is responsible for what is in your chart.
[Protecting your choice]	*Cover how and how often patients are informed; that they can decline at any point; that their care stays the same either way; that alternatives are available if they decline; and that communication comes in their preferred language on file.* **Example:** Before each visit where the tool is in use, you will be asked if you are comfortable with it. You can decline, your care stays the same either way, and other options are available. You can change your answer at any future visit.
[How to reach us with questions or concerns]	*Name one specific path: a person or a process. Keep it short.* **Example:** If you have questions, ask your provider or the front desk, or contact [named path].

As a best practice for any tool that operates during a clinical visit, inform the patient before each visit, confirm they understand what the tool does, and give them a genuine choice about its use.

> **LEGAL NOTE**
> Patient consent obligations vary by tool and setting. Informed-consent practices and any written disclosures must be reviewed by your legal and compliance team before use. This template is a communication guide only.

Part 3: Common Failure Modes

Talking points lose credibility in predictable ways. Use this list to check yours before they go out.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Watch for	Instead
Talking points written once and left unchanged	Refresh whenever an AI decision changes, a test concludes, or a new signal emerges.
Distributed before being stress-tested	Run the five-question review with a mixed group before every distribution.
Promising more than the health center can measure or deliver	State what you are watching and how you will know it is working.
Leading with the technology	Lead with the problem being solved and what changes for patients or staff.
Describing AI in terms that put technology ahead of mission	Use the terms your audience trusts: mission, patient access, financial sustainability, and capacity.
Treating values-based concerns as resistance to manage	Treat them as governance signals, name a path, and close the loop.
Naming a path for concerns that leads nowhere	Name a specific person or process, and confirm it is active first.
Inconsistent messaging across audiences	Board, staff, patients, and partners hear the same position.

KEY TAKEAWAYS BY ROLE

Leadership

- You approve talking points because they represent where your health center stands on AI.
- Lead every message with the why, drawn from the League's AI strategy: access, time for care, and financial sustainability.
- Refresh talking points whenever an AI decision changes, and use the worked examples as starting points.

AI Governance Hub members

- Run the four-step process for every AI decision: develop, stress-test, approve, then distribute.
- Stress-test with people outside the AI Governance Hub, including an IT or security reviewer when a message touches data or privacy.
- Track situations that still need talking points and bring them to the next meeting.

IT & Security

- When a message makes claims about data handling, privacy, or security, review it before it goes out.
- Confirm the Approved-Use Inventory Checklist is current so messages about current AI use stay accurate.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

This tool stays on the shelf at this stage. Your work is to understand the signal first, in Phase 1 — Listen and Learn (Section 3), before any tool decision.

→ Evaluating a specific AI tool?

Once your health center makes a governance decision, use Part 1 to build, test, and approve talking points, then the templates in Part 2. Go next to Phase 3 — Plan and Implement (Section 5).

→ Formalizing an existing AI tool?

Use the staff and external partner templates to communicate how the tool is entering your governance process. Go next to Phase 2 — Assess and Decide (Section 4) to begin the formal review.

→ Reassessing a changing AI tool?

Communicate the change to your board, staff, and relevant partners before it takes effect. Go next to Phase 2 — Assess and Decide (Section 4) to confirm whether the original approval still holds.

→ Checking an approved AI tool?

Use the board, staff, and patient templates to share what a performance or impact check found and what happens next. Go next to Phase 4 — Steward (Section 6).

→ Replacing or exiting an AI tool?

Communicate the decision to all four audiences before exit begins. If you are replacing the tool, return to Phase 2 — Assess and Decide (Section 4) for the incoming tool.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 3: Phase 1 — Listen and Learn

GOAL OF THIS SECTION

Build a standing listening practice that captures every AI signal, routes it to the AI Governance Hub, and keeps your health center current.

WHAT'S COVERED IN THIS SECTION

3.1 What "Listen and Learn" Means
Build a standing signal-and-literacy practice

3.2 Signal Types
Recognize the four kinds of AI signals

3.3 Patient, Staff, and Board Voice
Hear every level of the organization

3.4 Monitoring Regulation, Vendors, and Payers
Track outside changes on a schedule

TOOLS YOU'LL USE

AI Signal Intake Form Template
How do we capture and route an AI signal?

Patient AI Feedback Form Template
How do we hear from patients about AI in their care?

Regulatory and Compliance Monitoring Checklist
What do we watch for, and how often?

3.1 What "Listen and Learn" Means

This section is for the AI Governance Hub (the team your health center forms to oversee AI decisions), especially the AI Governance Hub Lead, the frontline staff representative, and whoever holds vendor and regulatory monitoring. Set up your AI Governance Hub and a regular meeting schedule (see Section 2 — Getting Started, subsection 2.4: The AI Governance Hub) before you begin here.

Your health center already gathers information about AI: vendor calls, staff questions, peer conversations, and features that appear in tools you already use. A signal is any indication that something AI-related needs attention. What makes this a system is consistency. Your AI Governance Hub catches signals on a regular schedule, from every part of your organization, through a process anyone can use. A regulatory change caught through routine monitoring can be absorbed on your own timeline, and a new feature in your health record system can be reviewed when your Hub chooses. A listening system needs two things: consistency and a low barrier to submit. When submitting a signal takes five minutes or less, people use it. The SEEDS of Innovation™ decision method that structures how your health center evaluates each signal is introduced in Section 1 — How This Playbook Works, subsection 1.3.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Building AI literacy runs alongside listening. Many people first meet AI in daily life: a voice assistant, a chatbot, a search engine that answers in sentences. That everyday experience is a starting point, and the questions staff bring from it are worth having at work. Three levels cover most needs:

- **Awareness (all staff):** what AI is, and where it already shows up in the work.
- **Practical use (staff who use AI tools daily):** how to use a tool safely, and what to do when an output looks wrong.
- **Decision-making (the AI Governance Hub and leadership):** how to evaluate a vendor, set policy, and monitor for harm.

Free, community-health-relevant training exists today, including courses built for safety-net settings and community health workers. See Appendix E: AI Literacy Resources for current options.

3.2 Signal Types

Signals come from inside and outside your organization, and each kind calls for a different response. Four categories cover most of what health centers encounter, and the AI Signal Intake Form Template captures all four. Any AI Governance Hub member or staff member can submit one.

Benefit or opportunity. These signals point to tools or features that appear to deliver value: restoring clinician time (such as ambient clinical documentation tools, which listen to and summarize a visit), supporting staff capacity (such as AI-assisted call center tools), or improving access (such as language support tools that reach patients in their primary language). A benefit confirmed at one health center still needs its own review against your conditions, which happens in Phase 2 — Assess and Decide (Section 4).

Harm, risk, or performance concern. These signals point to tools that add burden, create compliance exposure, or fall short of what was promised: a tool that makes more work instead of less, training that competes with patient-care time, patient trust affected by how AI changes the visit, or documentation that raises questions about accuracy. Some concerns appear only with specific groups. A documentation tool that works well in primary care may perform differently in behavioral health, and ambient tools can leave out the social-drivers-of-health documentation that matters for whole-person care. These signals call for an AI Governance Hub response and move into assessment.

Policy, regulatory, or vendor change. These signals come from outside and shape what happens inside: federal and state rules, vendor features turned on through a standard update, and payer policies affecting whether AI-assisted documentation is accepted. Massachusetts carries a specific restriction: MassHealth's Quality and Equity Incentive Program (QEIP) prohibits the use of AI for medical interpretation, which affects every health center serving patients through MassHealth. Subsection 3.4 covers the monitoring practice for this category.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI tool already in use without review. These signals surface when your health center finds AI already in place: a staff member built a process around a tool the Hub has yet to review, or a vendor feature activated months ago and kept running. When a tool is already in use, the goal is to bring it into governance: confirm what it does, check how it performs, and hold it to the same standard as any newly evaluated tool. These enter through the AI Signal Intake Form like any other signal.

> **IN PLAIN TERMS**
>
> Shadow AI means staff using AI tools the AI Governance Hub has yet to review or approve. It often enters through tools your health center already trusts, so standard review can miss it. The AI Signal Intake Form includes Track 2 so staff can surface these concerns at any time.

3.3 Patient, Staff, and Board Voice

Your governance process grows stronger when it hears from patients, frontline staff, and board members. Each sees AI from a different position, and their input is governance information: what your AI Governance Hub uses to make decisions.

Patient voice. Patient input matters most for AI tools that affect how care is communicated, how data is used, how consent is obtained, or how the visit feels. Patients often see what providers miss: whether an AI interaction felt trustworthy, whether they understood what was happening with their information, or whether the technology created a new barrier. Many patients served by community health centers face real barriers to digital participation: limited device access, low data plans, variable digital literacy, and real concerns about data collection. Approaches that work best include brief verbal check-ins during visits, paper forms at the front desk, engagement with patient advisory committee members on AI-specific questions, and listening in each patient's own language. The Patient AI Feedback Form Template gives your Hub a ready structure for this. When patterns surface across completed forms, the AI Governance Hub Lead logs them as signals through the AI Signal Intake Form and routes them for assessment.

> **IMPORTANT**
>
> This Playbook was built to bring patient voice into every AI decision your health center makes. When your AI Governance Hub reviews an AI tool, patient input is part of that review.

Frontline staff voice. Frontline staff (nurses, medical assistants, patient navigators, front-desk staff, care coordinators, and behavioral health workers) are the first to know when a tool adds work instead of removing it, the first to find workarounds, and the first to notice when something fails for a specific patient group. Talk with frontline staff before you test a new tool, include them while you set it up, and keep asking for their input after it is running. Their experience in the first 90 days is some of the most

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

useful information your Hub will receive. The AI Signal Intake Form includes Track 2 so staff can raise observations, questions, or concerns at any time.

Board voice. The board belongs in the AI governance conversation at a regular interval. What it needs is a periodic, plain summary: what AI tools are in use, what decisions have been made, what risks are being watched, and what the organization has learned. A brief AI update at a regular board meeting, or a short written summary beforehand, keeps the board informed while using one shared process. Board members who are also patients may surface patient-perspective signals, and those belong in the intake log like any other.

3.4 Monitoring Regulation, Vendors, and Payers

Some signals require you to go looking. The regulatory environment, vendor landscape, and payer (the insurance company or program that pays for care) policies affecting AI move on timelines outside your health center's control. Monitoring is how your organization stays current before something changes without warning. Even a small health center can do this, with the AI Governance Hub Lead holding monitoring directly. All it takes is a defined schedule and a consistent place to log what you find. The Regulatory and Compliance Monitoring Checklist provides that structure, in two parts: Part A for external monitoring and Part B for internal governance checks.

Federal and state policy. Federal guidance on AI in clinical settings, privacy updates, and documentation requirements all shape what your health center can do and must be able to show. Massachusetts adds state-level requirements, including data-privacy protections, Division of Insurance requirements, and Section 1557 of the Affordable Care Act (nondiscrimination requirements for health programs receiving federal funds). Massachusetts also carries the MassHealth QEIP restriction on AI for medical interpretation. Monitoring means knowing where to look and having someone check on a defined schedule. See Appendix C and Appendix D for policy and compliance detail.

Contracts and payer changes. Business Associate Agreements (BAAs) are contracts with companies that handle patient data on your behalf. Many were written before AI became part of these tools and may need updating to reflect what the tools now do with patient information. Monitoring includes checking whether your current agreements still cover your current tools; Phase 2 — Assess and Decide (Section 4), subsection 4.1 covers contract terms in more detail. Payer rules also vary: some accept AI-assisted documentation, and others are developing guidance on when AI use affects reimbursement. The MLCHC network is an active channel for sharing payer intelligence, where what one health center learns helps the others.

Vendor, health record system, and network updates. Your health record system (the electronic health record, or EHR) may receive AI features through a standard update, with no decision from your health center. Some EHR companies add new AI features and turn them on for everyone within a few months. Check your system's update notes before each update window, and review vendor communications and

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

contract terms on a regular schedule. The Massachusetts League of Community Health Centers (MLCHC) is an active voice on AI governance and shares intelligence through forums and advisories. When it does, that information belongs in your signal log. See Appendix D for current MLCHC contacts and channels.

CASE EXAMPLE: BAYSIDE HEARS THREE SIGNALS ABOUT THE SAME NEED

This example is hypothetical and depicts a composite community health center. It continues the Bayside story from Section 2.

The story so far. In Section 2, Bayside Community Health Center found two unreviewed AI tools in its building, built an inventory and a policy, and stood up its **AI Governance Hub**, the small group that makes its AI decisions. The AI Governance Hub paused a free transcription tool and sent the health record system's note-taking feature to full review.

The situation. Bayside's AI Governance Hub now meets monthly, and the AI Governance Hub Lead keeps the signal log. In one month, three things happen. A vendor shows an ambient documentation tool at a conference, a tool that listens during the visit and drafts the clinical note, and two physicians ask about it. A patient tells a medical assistant she felt unsure when a computer listened during her visit. And the monthly monitoring check finds a payer bulletin about AI-assisted documentation.

Each of these is a signal: any indication, from a person, a system, a vendor, a regulator, or a peer organization, that something AI-related needs attention. Here is how the AI Governance Hub catches all three.

Step 1 — AI Signal Intake Form Template *(How do we capture and route an AI signal?)*

The AI Governance Hub Lead fills out one intake form for each signal. The physicians' interest goes in through Track 1 as a "Benefit or opportunity," source "A company offering an AI tool reached out." The patient's comment goes in through Track 1 by the medical assistant who heard it, source "Something came up in a conversation with a patient or family member," category "Harm, risk, or performance concern."

Both are marked Routine priority, and the routing section says both hold for the next AI Governance Hub meeting, because they clearly belong together.

Writing signals down is the whole point of this step. A signal in the log can be acted on. A signal that stays in someone's head is lost.

Step 2 — Patient AI Feedback Form Template *(How do we hear from patients about AI in their care?)*

The note-taking feature from Section 2 is still under review, and some visits used it before the review began. After those visits, staff offer patients the short feedback form, on paper and in Portuguese.

The answers show a pattern. Asked "Did someone tell you that AI was used in your care today?", several patients answer "Not sure." Asked "How do you feel about AI being used in your care?", most answer "I have questions." Several want to know what happens to the recording of their visit. And on the question "Did your visit feel different because of AI?", a few patients write something positive: their doctor looked at them during the visit instead of typing.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

The AI Governance Hub Lead logs the whole pattern as one new signal through the AI Signal Intake Form Template. Patient reactions that were easy to miss have become something the AI Governance Hub can see and act on.

Step 3 — Regulatory and Compliance Monitoring Checklist *(What do we watch for, and how often?)*

Once a month, the AI Governance Hub Lead works through this checklist. Under "Contracts and Payer Changes" in Part A, the payer bulletin gets logged, with the date checked and a note in the finding-and-action field to confirm the payer's position before any pilot.

Part B checks the health center's own housekeeping: the signal log review confirms all three signals from this month are logged, categorized, and routed. Nothing is stuck.

The outcome. At its next meeting, the AI Governance Hub looks at the three signals side by side and sees them point at one need. Staff feel documentation burden. Patients have questions about recording. A payer is watching the same kind of tool.

The AI Governance Hub confirms the need is real, then routes the bundle to the AI Signal Decision Assessment, the full evaluation in Phase 2 — Assess and Decide (Section 4). Vendor contact waits until that assessment is done, and the patients' language questions travel forward with the signal.

What this teaches. The three tools work as a chain. The intake form catches what staff and patients notice. The feedback form turns scattered patient reactions into a clear pattern. The monitoring checklist catches what is changing outside the building and confirms nothing was missed. Because Bayside listened in all three directions before evaluating anything, Phase 2 starts with the full picture. Bayside's story continues at the end of Section 4.

For more detail, see Appendix C: Glossary, Appendix D: Tools, Templates and Resources, and Appendix E: AI Literacy Resources.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI Signal Intake Form Template

This tool answers: How do we capture and route an AI signal?

Phase 1 — Listen and Learn (Section 3)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A two-track form for capturing any AI signal. Track 1 takes any source; Track 2 takes staff concerns.	The submitter completes the relevant track and sends it to the AI Governance Hub Lead, on paper or by email.	The AI Governance Hub Lead logs, categorizes, and routes every signal, then hands it to the next phase for review.

Use **Track 1** for any AI-related signal from any source. Use **Track 2** to raise a staff concern, with the option to stay anonymous. Complete the track that fits, then send it to the AI Governance Hub Lead. Have a brief note on what you observed ready; a sentence or two is enough.

Track 1 · General signal intake

For any staff member, AI Governance Hub member, or leader

Date submitted	Submitted by (name and role)

Where did this signal come from?

☐ A company offering an AI tool reached out (by phone, email, or in person)

☐ A new feature appeared in a system the health center already uses

☐ A staff member mentioned something (with curiosity, a question, or a concern)

☐ Something came up in a conversation with a patient or family member

☐ A peer organization or the Massachusetts League of Community Health Centers (MLCHC) network shared something

☐ A regulatory, payer, or policy update

☐ Other (describe below)

Priority:

☐ Urgent: needs AI Governance Hub review before the next scheduled meeting

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

☐ Routine: add to the signal log and review at the next AI Governance Hub meeting

☐ Unsure: AI Governance Hub Lead to determine

Describe the signal

What happened, and why it may need attention

Supporting materials

Attach, or note where to find them

Track 2 · Staff concern

Raise a concern, with the option to stay anonymous

Your voice matters

This form is a safe space for raising concerns. [Health Center Name] keeps Track 2 submissions separate from performance reviews and disciplinary decisions. You may submit anonymously. If you choose to share your name, the AI Governance Hub Lead holds your identity in strict confidence, with the same care as any governance record.

A note on data privacy: if you realize a submission includes Protected Health Information (PHI) or patient data by mistake, contact IT right away and hold that submission until IT advises next steps.

Date submitted	Submitted by (optional: you may remain anonymous)

What are you concerned about?

☐ An AI tool being used at this health center that was never reviewed or approved

☐ An AI tool that is creating more work, causing problems, or performing below expectations

☐ How AI is affecting patients (something I witnessed or heard directly)

☐ Something I am unsure how to raise through normal channels

☐ Other (describe below)

Has this concern been raised with anyone else at the health center?

☐ Yes ☐ No ☐ I am not sure

Follow-up:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

☐ I would like the AI Governance Hub Lead to follow up with me. Contact:

_____

☐ I want this documented. Follow-up is optional for me.

Describe what you observed or are concerned about

What you saw or heard, and why it concerns you

Thank you for helping to improve care and workflows at your health center.

For the AI Governance Hub Lead · Signal routing

If you are submitting this form, you are done. Stop here.

This section is completed by the AI Governance Hub Lead after receiving a submission.

Date received	Completed by (AI Governance Hub Lead)

Added to signal log:

☐ Yes, date: _____ ☐ Pending, reason: _____

Signal category (see Section 3.2: Signal Types):

☐ Benefit or opportunity

☐ Harm, risk, or performance concern

☐ Policy, regulatory, or vendor change

☐ AI tool already in use without review

Initial routing:

☐ Holds for next AI Governance Hub meeting (routine review)

☐ Moves to assessment now (urgent, or tool already in use without review)

☐ Needs more information before routing (follow-up required)

Routed to the AI Signal Decision Assessment (the full evaluation in Phase 2 — Assess and Decide, Section 4):

☐ Yes, date: _____

☐ Pending AI Governance Hub review

☐ Holds for now, reason: _____

AI Governance Hub Lead notes

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Routing rationale, follow-up owner, and any next step

KEY TAKEAWAYS BY ROLE

Leadership

- This form is the front door for every AI signal at your health center. Supporting its use shows staff that leadership takes their observations seriously.
- Review the signal log regularly with the AI Governance Hub Lead to stay aware of what is surfacing.

AI Governance Hub members

- Every signal submitted through this form is logged and reviewed by the AI Governance Hub Lead, who determines the next step.
- Track 2 submissions require confidentiality. Handle them with the same care as any governance record.

IT & Security

- System changes, new vendor features, and access requests are signals. When you see AI capabilities in tools your health center already uses, submit through Track 1.
- Track 1 submissions from IT & Security help the AI Governance Hub identify tools already in use without review.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

You identified a need or gap that AI might address. Submit through Track 1 to log the signal. The AI Governance Hub Lead reviews it and routes it to the AI Signal Decision Assessment in Phase 2 — Assess and Decide (Section 4) if it warrants full evaluation.

→ Evaluating a specific AI tool?

You have a named capability on the table. This form may already have captured the signal that started this path. If new signals surface during evaluation, submit them through Track 1. Your evaluation tools are in Phase 2 — Assess and Decide (Section 4).

→ Formalizing an existing AI tool?

A tool appeared in a workflow, a vendor activated a feature, or staff began using an AI tool on their own. Submit through Track 1 to log the tool, or Track 2 to raise it as a concern. The AI Governance Hub Lead routes it to the AI Signal Decision Assessment in Phase 2 — Assess and Decide (Section 4) for retroactive review.

→ Reassessing a changing AI tool?

The vendor updated features, expanded scope, or changed contract terms. Submit through Track 1 to log the change. The AI Governance Hub Lead decides between a fresh review in Phase 2 — Assess and Decide (Section 4) and a stewardship check in Phase 4 — Steward (Section 6).

→ Checking an approved AI tool?

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

A scheduled review is due, an incident occurred, or feedback raised a concern about a tool in use. Submit through Track 1 to log what was observed. The AI Governance Hub Lead routes it to the Stewardship Review Checklist in Phase 4 — Steward (Section 6).

→ Replacing or exiting an AI tool?

Your health center is exiting an AI tool or beginning a replacement. Submit through Track 1 to log the transition signal. The AI Governance Hub Lead routes it to Phase 4 — Steward (Section 6) to manage the exit, and back to Phase 2 — Assess and Decide (Section 4) for any incoming tool.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Patient AI Feedback Form Template

This tool answers: How do we hear from patients about AI in their care?

Phase 1 — Listen and Learn (Section 3)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A short, plain form asking patients about AI in their care: what they were told, how it felt, their questions.	Offer it after a visit that used AI, on paper, portal, or text. Staff bring forms to the Hub Lead.	The Hub Lead logs forms, tracks patterns on a schedule, and routes concerns through the AI Signal Intake Form.

Before You Use This Form

This form works for any AI tool that affects how patients experience their care, whether clinical (ambient documentation, AI-assisted call centers, language support tools) or administrative (scheduling, patient communication, or other AI-powered processes patients interact with). Adapt the language, format, and delivery method (paper, patient portal, text or email link) to fit your patient population, using the engagement approaches that already work for your community.

Before offering this form, confirm two things:

- Your health center has language support that matches your patient population, and a process for translating completed responses back into English so the AI Governance Hub can log and review them.
- Staff are trained to assist patients who need help completing it.

The “Need Help” box on the patient-facing section tells patients they can ask a team member for help. Make sure staff are ready to give that help.

Completed forms go to the AI Governance Hub Lead. This is governance information. It stays with the AI Governance Hub for evaluation and assessment, separate from patient satisfaction surveys, chart notes, and care records.

IMPORTANT

This Playbook was built to bring patient voice into every AI decision your health center makes. This form gives your AI Governance Hub a structured way to hear directly from patients about their experience with AI in their care.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

FORM TO GIVE TO THE PATIENT

NEED HELP WITH THIS FORM?

Ask a team member. They can help you fill it out.

If you prefer to complete this form in another language, please ask.

About AI and Your Care

This health center uses AI in some parts of your care.

AI stands for Artificial Intelligence. AI is a kind of computer technology. It learns from a lot of information. Then it uses what it learned to do tasks that usually need human thinking. Some examples: writing notes from your visit, answering questions, finding patterns in health records, or making suggestions to your care team.

This short form asks what you think. Your answers help us do better.

Filling this out is your choice. Your care stays the same either way.

Part 1: What You Were Told

1. Did someone tell you that AI was used in your care today?

☐ Yes ☐ No ☐ Not sure

2. Did you understand what the AI tool does?

☐ Yes, I understood

☐ Not really

☐ I want to know more

Part 2: How It Felt

3. How do you feel about AI being used in your care?

☐ I am okay with it

☐ I am not sure how I feel

☐ I do not like it

☐ I have questions

4. Did your visit feel different because of AI?

☐ Yes ☐ No ☐ Not sure

If yes, tell us what felt different:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Part 3: Your Questions and Concerns

5. Do you have a question or concern about AI in your care?

☐ Yes ☐ No

If yes, write it here:

6. Would you like someone to talk with you about this?

☐ Yes ☐ No

If yes, how can we reach you?

These last two lines are optional. Sharing them is your choice.

Date of your visit	
Your name (optional)	

Thank you for helping to improve care and workflows at [Health Center Name]!

KEY TAKEAWAYS BY ROLE

Leadership	AI Governance Hub members	IT & Security
• This form gives your health center a structured way to hear from patients about AI in their care. • Patient feedback patterns inform AI Governance Hub decisions and strengthen the Pillar 2: Everyone-Centered Design assessment. This pillar evaluates whether an AI tool centers the needs of everyone it affects, especially those with the highest barriers. • Confirm that translation resources are in place so the form reaches patients in their	• Establish a regular schedule for reviewing completed forms and identifying trends. • When patterns surface a concern, log it as a signal through the AI Signal Intake Form and route it for assessment.	• Confirm that digital delivery channels (patient portal, text or email links) are set up and delivering forms to patients and receiving responses. • Verify that completed form responses route to the AI Governance Hub Lead and remain separate from clinical records and the patient's chart. • Confirm that translation workflows are in place: forms delivered in the patient's preferred language, and completed responses translated back into English for the AI

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

preferred language and completed responses return to the AI Governance Hub in English.		Governance Hub to log and review.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**

Set up this form before any AI tool goes live so you are ready to hear from patients from day one. Preparation supports Phase 1 — Listen and Learn (Section 3).

→ **Evaluating a specific AI tool?**

Collect patient input during a pilot to inform the AI Governance Hub's assessment. Feed results into Phase 2 — Assess and Decide (Section 4).

→ **Formalizing an existing AI tool?**

Use this form to hear from patients already experiencing a tool that arrived without a formal review. Bring patterns to the AI Governance Hub and enter Phase 1 — Listen and Learn (Section 3).

→ **Reassessing a changing AI tool?**

Hear whether patients notice the change and how it affects their experience. Route findings to Phase 2 — Assess and Decide (Section 4) to determine whether the change requires a fresh assessment.

→ **Checking an approved AI tool?**

Patient feedback patterns are a key input for stewardship reviews. Route findings to Phase 4 — Steward (Section 6) through the AI Signal Intake Form.

→ **Replacing or exiting an AI tool?**

Capture patient experience with the outgoing tool before removal. Findings inform the exit process and, if replacing, feed into Phase 1 — Listen and Learn (Section 3) for the incoming tool.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Regulatory and Compliance Monitoring Checklist

This tool answers: What do we watch for, and how often?

Phase 1 — Listen and Learn (Section 3)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A checklist for tracking outside changes and internal governance checks that affect your AI tools.	Run it monthly, or quarterly at minimum. Check each source, log the date, note findings.	The AI Governance Hub Lead runs it; the full Hub reviews findings at its regular meeting.

How to Use This Checklist

The rules and tools around AI change often. A policy update, a vendor feature release, or a payer requirement can affect patient privacy, reimbursement, or a tool your health center already uses. This checklist keeps your health center current so every change gets the attention it requires.

The AI Governance Hub Lead runs this checklist, collects findings, and brings them to the AI Governance Hub for review. Any staff member can complete individual items. For each item, check the source listed, log the date, and flag anything that needs attention.

At larger health centers, specific items may be assigned to IT, compliance, or operations staff. The Hub Lead collects their reports and completes the log.

> **IN PLAIN TERMS**
> A Business Associate Agreement (BAA) is the data-handling contract with any company that manages patient information on your behalf.

How Often to Run It

Set a schedule that matches your health center's pace.

- Monthly (default) — recommended for active governance. Run at each AI Governance Hub meeting cycle.
- Quarterly (minimum) — the floor for any health center. Check every monitoring item at least once per quarter.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- More frequent — health centers with many active AI tools, or during periods of rapid policy change, may check some items more often.

Checklist Run Header

Complete the header below once per run. For each checklist item, log the date you checked it and any findings or actions.

Period covered	
Completed by	
Date reviewed at AI Governance Hub meeting	

Part A: External Monitoring

These items track changes in the policy, regulatory, vendor, and payer environment that may affect how your health center governs and uses AI tools.

Part A covers two purposes. Some items ask you to watch what is coming: new rules, proposed legislation, or policy direction that may affect your health center later. Other items ask you to check what applies now: whether your current agreements, vendor settings, and payer requirements still match how your AI tools operate today. Each item is labeled so you know which purpose it serves.

FEDERAL AND STATE POLICY		
QUARTERLY	☐ **Federal AI and health IT policy — what is coming** **Where to look:** U.S. Department of Health and Human Services (HHS), Centers for Medicare and Medicaid Services (CMS), and Office of the National Coordinator for Health IT (ONC) public guidance and rule-making. Focus on proposed rules, executive orders, and guidance affecting clinical AI, documentation, patient privacy, and reimbursement. See Appendix D for a curated list of sources.	Date checked: Finding / action:
QUARTERLY	☐ **Federal AI and health IT policy — what applies now** **Where to look:** HHS Office for Civil Rights guidance on the Health Insurance Portability and Accountability Act (HIPAA) and the Health Information Technology for Economic and Clinical Health Act (HITECH). Review whether any current requirement has changed in a way that affects how your health center handles patient data with AI tools. Check MLCHC network advisories or your legal counsel for state-level privacy updates. See Appendix D.	Date checked: Finding / action:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUARTERLY	☐ **Massachusetts state policy channels** **Where to look:** Massachusetts Health Policy Commission, Division of Insurance, and Executive Office of Health and Human Services (EOHHS) policy communications. Review for new state-level requirements affecting AI use at your health center. See Appendix D.	Date checked: Finding / action:
QUARTERLY	☐ **MassHealth restriction on AI for medical interpretation** **Where to look:** MassHealth Quality and Equity Incentive Program (QEIP) requirements and any updates to the restriction on AI use for medical interpretation. This restriction affects every health center serving patients through MassHealth. Check for updates or changes.	Date checked: Finding / action:
CONTRACTS AND PAYER CHANGES		
QUARTERLY	☐ **Business Associate Agreements for active AI vendors** **Where to look:** Review the Business Associate Agreement (BAA) for every active AI vendor. Confirm it still accurately covers what the tool currently does with patient data.	Date checked: Finding / action:
QUARTERLY	☐ **Payer policy on AI-assisted documentation and reimbursement** **Where to look:** Each major payer's provider bulletins, policy updates, and manual changes, plus payer intelligence shared through MLCHC network communications. Focus on whether AI use affects what any payer will accept or reimburse. See Appendix D.	Date checked: Finding / action:
VENDOR, HEALTH RECORD SYSTEM, AND NETWORK UPDATES		
MONTHLY	☐ **Vendor communications and feature announcements** **Where to look:** Release notes, product update emails, and account communications from each active AI vendor. New features may activate automatically; check before each update window. See Appendix D.	Date checked: Finding / action:
MONTHLY	☐ **Health record system AI features and updates** **Where to look:** Electronic health record (EHR) vendor release notes and system update announcements. AI features may be turned on by default during standard updates. Confirm whether any new AI capability has been activated since the last check.	Date checked: Finding / action:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

MONTHLY	☐ **MLCHC IT Forum** **Where to look:** The MLCHC IT Forum communications, an opt-in network channel carrying AI governance intelligence, payer updates, and regulatory advisories for Massachusetts community health centers. To receive these, sign up through MLCHC; if your health center is missing them, contact MLCHC.	Date checked: Finding / action:

Part B: Internal Governance

These items check whether your health center's own governance process is current and working as designed. They use your own records and active tools only.

IS YOUR SIGNAL LOG AND APPROVED-USE LIST CURRENT?		
MONTHLY	☐ **Signal log review** **Where to look:** The AI Signal Intake Form log. Confirm all signals submitted since the last AI Governance Hub meeting have been logged, categorized, and routed. Flag any signal still waiting for a routing decision.	Date checked: Finding / action:
QUARTERLY	☐ **Approved-use list check** **Where to look:** The Approved-Use Inventory Checklist. Confirm the list reflects all AI tools currently in use. Flag any tool in active use that is missing from the list.	Date checked: Finding / action:
ARE TOOLS BEING USED THE WAY THEY WERE APPROVED?		
QUARTERLY	☐ **AI policy compliance check** **Where to look:** The organization's AI policy and the Approved-Use Inventory Checklist. For each active AI tool, confirm it is operating within the scope defined by the policy and the approved-use list. Look for a tool used in a way that was never reviewed, a tool whose use has expanded beyond its original approval, or a tool that should have been decommissioned but is still running.	Date checked: Finding / action:
PLAYBOOK AND NETWORK UPDATES		
QUARTERLY	☐ **Playbook update triggers** **Where to look:** Findings from this checklist and AI Governance Hub meeting records. If any external change requires an update to your governance process, flag it and bring it to the AI Governance Hub. The primary maintenance process lives in Section 7: Keeping the System Running.	Date checked: Finding / action:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUARTERLY	☐ **Network knowledge transfer** **Where to look:** AI Governance Hub meeting records and completed checklist logs. Document anything learned this quarter worth sharing with MLCHC or the broader network, and share a summary with health center leadership. Use the Case Example Sharing Template (Section 6) to structure the submission.	Date checked: Finding / action:

KEY TAKEAWAYS BY ROLE

Leadership

- Use this checklist to confirm your health center is actively tracking the outside changes that affect AI decisions.
- Review completed checklists quarterly, via report-out from the AI Governance Hub, to stay ahead of regulatory, vendor, and payer shifts.

AI Governance Hub members

- Completed checklists feed your regular meeting. Review findings and route anything requiring action through the AI Signal Intake Form.
- Assign monitoring items across the team so the workload is shared and every source gets checked.

IT & Security

- Own the vendor and health record system items. You are closest to release notes, feature activations, and update windows.
- Flag any AI capability that activated without a governance decision, and route it through the AI Signal Intake Form (Section 3).

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**
A monitoring finding surfaced a possible need. Route it through the AI Signal Intake Form, then move to Phase 2 — Assess and Decide (Section 4).

→ **Evaluating a specific AI tool?**
Confirm federal, state, and payer requirements are current, then continue in Phase 2 — Assess and Decide (Section 4).

→ **Formalizing an existing AI tool?**
This checklist flagged a tool in use without a formal review. Submit it through the AI Signal Intake Form and begin Phase 2 — Assess and Decide (Section 4).

→ **Reassessing a changing AI tool?**
A check revealed a contract, payer policy, or requirement shifted for a tool in use. Return to Phase 2 — Assess and Decide (Section 4) to confirm the approval still holds.

→ **Checking an approved AI tool?**
A scheduled review is due. Confirm the environment has stayed stable and pair findings with the Stewardship Review Checklist in Phase 4 — Steward (Section 6).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

→ **Replacing or exiting an AI tool?**

A regulatory change, vendor shutdown, or failed review means a tool must be retired. Document the external conditions, then move to Phase 4 — Steward (Section 6).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 4: Phase 2 — Assess and Decide

GOAL OF THIS SECTION

Work every AI signal through all five SEEDS pillars, with evidence, to reach a documented recommendation and confirm readiness before go-live.

WHAT'S COVERED IN THIS SECTION

4.1 Vendor Evaluation
Capture vendor facts before deciding.

4.2 Applying SEEDS to Signals
Structure the Hub's reasoning across five pillars.

4.3 Maturity Check
Reveal how deeply the Hub reasoned.

4.4 Behavioral Health, SDoH, and Sensitive Populations
Add care for higher-stakes signals.

4.5 Pre-Deployment Impact Checklist
Confirm readiness before go-live.

TOOLS YOU'LL USE

Vendor Evaluation Checklist
What do we need to know before and during vendor evaluation?

AI Signal Decision Assessment
Should we move forward, and how strong is our reasoning?

Pre-Deployment Impact Checklist
Is this AI tool safe to deploy?

4.1 Vendor Evaluation

Phase 2: Assess and Decide begins once a signal has been documented in the AI Signal Intake Form Template and routed from Phase 1: Listen and Learn. This work is for the AI Governance Hub, especially the Hub Lead, clinical leadership, and the frontline staff representative.

IN PLAIN TERMS

A signal is any indication that something AI-related needs your AI Governance Hub's attention, such as a vendor pitch, a staff observation, a patient concern, a payer change, or AI that arrived without a decision.

When a signal involves a new vendor (the company providing the tool), the AI Governance Hub begins with the Vendor Evaluation Checklist, before any conversation with that vendor takes place. When a

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

health center expands an existing tool to a new use, there is often no new vendor to evaluate, and the Hub moves directly to the AI Signal Decision Assessment (Section 4.2).

The Vendor Evaluation Checklist runs in two parts, both organized by the five SEEDS pillars. Part 1, Pre-Meeting Intake, is where the Hub Lead assigns a small team, typically the IT or data lead plus legal or compliance counsel, to research the vendor from publicly available materials and send written questions before the first meeting. It is a living document, updated as answers come in. Part 2, Evaluation Rubric, is where the team judges whether the vendor adequately addressed each item during and after the meeting, marking Yes, Partially, or Not yet, and records the evidence.

The checklist also asks whether the vendor has experience with Federally Qualified Health Centers (FQHCs) and community health centers, offers grant-eligible or phased pricing, and can provide references from safety-net settings similar to yours. The checklist captures facts. Evaluation of fit happens in the AI Signal Decision Assessment.

The Commonwealth Purchasing Group (CPG), a Massachusetts League of Community Health Centers company, negotiates discounted rates with vendors for its members, and membership is free. Before or during vendor evaluation, health centers can check the CPG vendor catalog at savewithcpg.com to see whether a solution that addresses the need is already available at a discounted rate. A completed checklist also strengthens a health center's position in collective negotiation by documenting the specific questions a vendor must answer.

4.2 Applying SEEDS to Signals

The AI Signal Decision Assessment is the primary instrument for Phase 2, and every documented signal moves through it. The AI Governance Hub works through all five SEEDS of Innovation™ pillars in sequence, introduced in Section 1.3: The Decision Method: SEEDS of Innovation™. Each pillar examines a different dimension of the signal: the problem being solved, the people affected, the fit with how care is actually delivered, the data conditions, and who is accountable over time.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Figure 4.1. How the five-pillar assessment produces the AI Governance Hub's recommendation

The assessment works in three steps. First, within each pillar the AI Governance Hub answers Yes or No to each question, with documented evidence; the questions are organized across the five reasoning-maturity levels described in Section 4.3. Second, the AI Governance Hub records the earned reasoning level for that pillar, which is the highest level where every question at that level, and every level below it, is answered Yes. Third, when all five pillars are complete, the Hub reviews the reasoning profile across the pillars, where the lowest pillar matters most, and reaches a recommendation: Approve, Approve with conditions, Hold, or Do not proceed.

Because the Playbook assumes a buy-first approach, an Approve or Approve-with-conditions recommendation points to the Vendor Evaluation Checklist to compare market options before any decision to build. When a specific vendor is already at the table, the order reverses: the Vendor Evaluation Checklist runs first, and its findings feed this assessment (see Section 4.1). The AI Governance

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Hub documents the recommendation in the Design Plan Template (Section 5: Phase 3 — Plan and Implement).

4.3 Maturity Check

The AI Signal Decision Assessment includes a maturity check built into each pillar. It is a reflection tool, a way for the AI Governance Hub to examine the quality of its own reasoning.

Inside each pillar, the assessment asks questions across five reasoning-maturity levels: Reflexive, Protective, Normative, Integrative, and Regenerative. The AI Governance Hub answers each question Yes or No, with evidence. The earned level for each pillar is the highest level where every question at that level, and every level below it, received a Yes answer.

When all five pillars are complete, a reasoning summary shows the earned level across all five pillars side by side. This summary is a self-awareness tool. An AI Governance Hub working at the first level on a given signal is at an early stage of assessing that signal. The maturity check makes that visible, so the Hub can name what would need to be true to reason at the next level. The purpose is growth over time. The maturity model is introduced in Section 1.3 and detailed in Appendix A: SEEDS of Innovation™.

> **IMPORTANT**
> A person stays in charge at every step. This Playbook uses a human-in-the-loop approach throughout: AI tools inform the AI Governance Hub's reasoning, but people make every recommendation and every go-live decision.

4.4 Behavioral Health, SDoH, and Sensitive Populations

Signals that involve behavioral health care, substance use disorder treatment, reproductive health, social drivers of health (SDoH) data, or other sensitive clinical areas call for additional attention in Pillar 2: Everyone-Centered Design of the AI Signal Decision Assessment.

The additional attention recognizes that the stakes for getting it wrong are higher. Data requirements are more specific, and the regulatory overlay is more complex. 42 CFR Part 2, a federal law that gives records about substance use treatment stronger privacy protections than standard HIPAA (the federal law that protects patient health information), applies to substance use records and imposes requirements beyond standard privacy rules. The patient population may be more vulnerable to harm if the tool fails or if data is mishandled.

> **LEGAL NOTE**

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> Before completing Pillar 2: Everyone-Centered Design for any signal in these categories, involve legal or compliance counsel, either in the session or through a consultation beforehand. If counsel is unavailable, pause Pillar 2 on that signal and schedule the consultation before proceeding.

The AI Signal Decision Assessment includes specific prompts for these signal types. See Appendix C: Glossary for definitions of the regulations referenced here, including 42 CFR Part 2.

4.5 Pre-Deployment Impact Checklist

Before any AI tool goes into regular use, the AI Governance Hub runs the Pre-Deployment Impact Checklist. This applies whether the tool is new, expanded to a new use, or returning after a significant change.

The checklist is organized by the five SEEDS pillars. For each question, the Hub answers Yes, Partially, or Not yet, with documented evidence such as a signed record, a completed test log, or a verified result. Across the pillars, the checklist confirms the tool's risk level and whether it touches patient health information (Systemic Signal Assessment); patient disclosure, accessibility, and bias checks (Everyone-Centered Design); regulatory authorization, integration testing, staff training, and override capability (Ecosystem Alignment); the signed Business Associate Agreement and verified encryption and logging (Data Sovereignty); and a rollback plan and monitoring cadence (Stewardship).

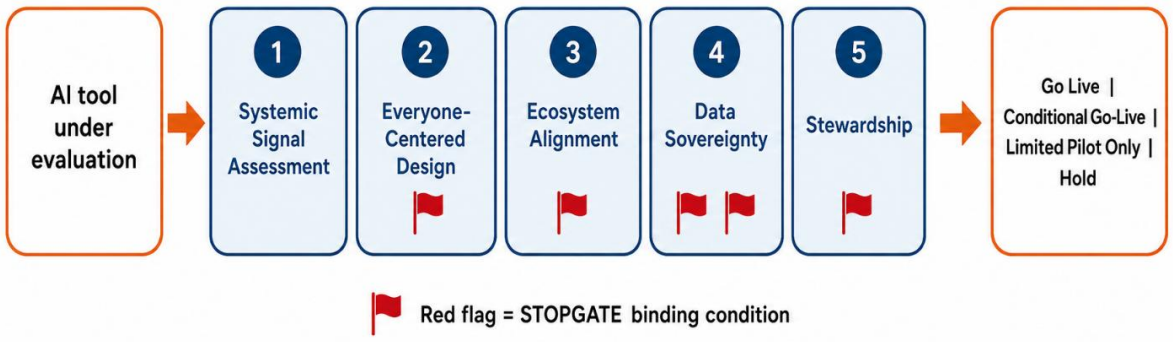

Figure 4.2. How an AI tool moves through the five SEEDS pillars to a go-live decision

Five questions are STOPGATE binding conditions, one owned by each sign-off role and distributed across the pillars: patient disclosure in the patient's primary language, regulatory authorization for the intended use, a signed BAA on file, encryption and logging verified in the health center's own

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

environment, and a documented rollback plan. A Partially or Not yet on any binding condition halts go-live until it is resolved with documented evidence, because going forward would create harm or exposure that cannot be undone.

When the checklist is complete, the AI Governance Hub selects one of four go-live decisions: Go Live (all questions Yes with evidence, all binding conditions cleared); Conditional Go-Live (all binding conditions cleared, with documented resolution plans, owners, and deadlines for any remaining Partially answers); Limited Pilot Only (one or more binding conditions still partially met, scope restricted to a defined pilot pending resolution); or Hold (unresolved binding conditions with an unclear resolution path). A Hold or Limited Pilot decision is a governance strength: the process caught a gap before it reached patients or staff. A completed Pre-Deployment Impact Checklist, with binding conditions cleared, is one of the formal conditions for moving forward to Section 5: Phase 3 — Plan and Implement.

CASE EXAMPLE: BAYSIDE DECIDES ABOUT THE DOCUMENTATION TOOL

This example is hypothetical and depicts a composite community health center. It continues the Bayside story from Section 3.

The story so far. Bayside Community Health Center stood up its **AI Governance Hub** in Section 2. In Section 3, the AI Governance Hub logged three signals that point at one need: physicians want an ambient documentation tool, a tool that listens during the visit and drafts the clinical note; patients have questions about being recorded; and a payer is watching the same kind of tool.

The situation. Everything Bayside has gathered now comes together: the note-taking feature sent to full review in Section 2, and the three signals from Section 3. The medical director is interested in the vendor's ambient documentation tool. The vendor has references from two urban health systems.

Interest is a starting point. Before anything goes live, the AI Governance Hub needs facts, a reasoned recommendation, and a safety check. Those are the three tools of this section, in order.

Step 1 — Vendor Evaluation Checklist *(What do we need to know before and during vendor evaluation?)*

This checklist gathers facts; the evaluation of fit comes in Step 2. Before meeting the vendor, the IT lead and compliance counsel fill out Part 1, Pre-Meeting Intake, with what they can learn on their own, and send the vendor written questions. In the meeting, they work through Part 2, the Evaluation Rubric, marking each item Yes, Partially, or Not yet with the evidence beside it.

The answers tell a clear story. Data ownership comes back Yes: Bayside owns its data. Accuracy in Portuguese comes back Not yet: every reference site works in English only. Safety-net experience comes back Partially: the vendor has no health-center reference to offer.

One more answer gets flagged as a STOPGATE, a binding condition that halts go-live until it is resolved: the vendor may use Bayside's data to train its AI models only with written authorization, and that authorization is still open. The STOPGATE travels forward with the findings.

Step 2 — AI Signal Decision Assessment *(Should we move forward, and how strong is our reasoning?)*

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

The AI Governance Hub takes the facts through the five SEEDS pillars. Each pillar earns a reasoning level, and the levels are cumulative: a level is earned only when every question at that level and below is answered Yes.

The problem itself is well understood: documentation burden shows up in staff surveys, and AI is a fitting answer (Pillar 1: Systemic Signal Assessment, Level 4: Integrative). The weak spots are just as clear. Portuguese accuracy is unproven (Pillar 2: Everyone-Centered Design, Level 2: Protective). Health-record integration is unreviewed and there is no peer reference (Pillar 3: Ecosystem Alignment, Level 2: Protective). The privacy agreement is in place, but the model-training authorization is unresolved (Pillar 4: Data Sovereignty, Level 3: Normative). No one yet owns ongoing monitoring (Pillar 5: Stewardship, Level 2: Protective).

The lowest pillar matters most, and three pillars sit at Level 2. From the assessment's four options, the AI Governance Hub selects “Approve with conditions.”

The main conditions are specific, and each one can be enforced. The vendor commits in writing, with a firm date, to deliver proven Portuguese-language accuracy and Portuguese patient-disclosure materials within six months. The tool becomes standard practice only when it serves Portuguese-speaking patients as well as English-speaking ones; a pilot for some patients is a step toward serving all, never the end point. If the vendor misses the date, the pilot ends. That last condition matters because AI vendors are new companies, and some of them close or change direction; a written date protects Bayside either way.

The approval also carries three supporting conditions, smaller requirements the health center completes alongside the vendor's commitment: the vendor's authorization to use Bayside's data for model training gets resolved in writing, the tool's connection to the health record system gets tested, and a named clinician reviews note accuracy every month. All of the conditions and unresolved gaps carry forward into the Design Plan Template.

Step 3 — Pre-Deployment Impact Checklist *(Is this AI tool safe to deploy?)*

This is the safety check right before launch, and every Yes must rest on documented evidence, because vendor assurances alone are insufficient. The AI Governance Hub answers each question Yes, Partially, or Not yet.

The signed privacy agreement, verified encryption and logging, and a rollback plan (the plan for switching the tool off safely) all clear. One binding condition comes back Partially: patient disclosure in the patient's own language, because the Portuguese materials are promised but have yet to arrive.

A STOPGATE at Partially means one decision from the checklist's four options: “Limited Pilot Only,” scope restricted to a defined pilot pending resolution. The pilot will be short, English-language, small in scope, and on a written path to serving every patient, with the vendor's dated commitment as the condition for continuing.

The outcome. The vendor signs the commitment. The decision is recorded: Approve with conditions, Limited Pilot Only, with a firm date and a clear exit. Turning this decision into a real launch is the work of Section 5.

What this teaches. The three tools work as a chain. The Vendor Evaluation Checklist captured the facts. The AI Signal Decision Assessment turned facts into a reasoned recommendation with conditions and a deadline. The Pre-Deployment Impact Checklist's safety check caught the disclosure gap before any patient was exposed to it. A pilot that reaches only some patients is acceptable only as a dated, written step toward reaching all of them. The deadline and the exit

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

are what keep any group of patients from being left behind. Bayside's story continues at the end of Section 5.

REFERENCE APPENDICES

- **Appendix A: SEEDS of Innovation™ —** Full framework, maturity table, and worked examples for the five-pillar decision method used throughout this section.
- **Appendix B: Implementation Approach Reference —** Framework descriptions and sourcing notes for the Playbook's readiness tools.
- **Appendix C: Glossary —** Plain-language definitions of terms and regulations referenced here, including 42 CFR Part 2 and HIPAA.
- **Appendix D: Tools, Templates and Resources —** Full index of the Playbook's tools and templates.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Vendor Evaluation Checklist

This tool answers: What do we need to know before and during vendor evaluation?

Phase 2 — Assess and Decide (Section 4)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A structured checklist that captures the facts about a vendor's AI product, before and during the first meeting.	Part 1 records pre-meeting research. Part 2 evaluates the vendor across the five SEEDS pillars.	The Hub Lead assigns a team to complete both parts. Findings feed the AI Signal Decision Assessment.

This is an information-gathering checklist. Use it to capture what the vendor says: the facts your AI Governance Hub needs to run the AI Signal Decision Assessment. Evaluation of fit happens later. Honest capture matters more than tidy answers.

How to use this checklist

The checklist runs in two parts, both organized by the five SEEDS pillars:

- Part 1, Pre-Meeting Intake: gather facts from your research and from written questions sent to the vendor in advance. Treat it as a living document, updated across more than one round of gathering.
- Part 2, Evaluation Rubric: judge whether the vendor adequately addressed each item, marking Yes, Partially, or Not yet, and record the evidence. You may complete it across more than one meeting.
- The AI Governance Hub Lead assigns a small team, typically the IT or data lead plus legal or compliance counsel, to prepare Part 1. The full team works through Part 2 during the meeting.
- Legal or compliance counsel reviews the Data Sovereignty questions before any contract discussion moves forward.
- Completed findings feed the AI Signal Decision Assessment, where SEEDS pillars structure the governance decision, and the Pre-Deployment Impact Checklist, which confirms readiness before go-live.

Reading the Scope Tags

Every question in this tool applies to the Buy scenario, when the health center is evaluating an AI tool from a vendor. Some questions also carry a tag:

Context — where the tool operates:

- **[Clinical]** applies when the AI tool supports direct patient care or clinical decisions.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- **[Admin]** applies when the AI tool supports administrative or operational functions.
- No context tag means the question applies in both settings.

Binding condition:

- **STOPGATE** halts the go-live process until it is resolved, because going forward would create harm or exposure that cannot be undone after the fact.

Bracketed source keys identify the framework each question draws on; full citations appear in Appendix B: Implementation Approach Reference.

Vendor Identification

Complete this section for every vendor evaluation. Use one checklist per vendor.

Vendor / Company Name	
Product Name	
Contact Name and Title	
Contact Email / Phone	
Date of Meeting	
AI Governance Hub Team Members Present	

Part 1: Pre-Meeting Intake

This is a living document. Fill it in from your research and from written questions sent to the vendor before the meeting, and expect to update it across more than one round of gathering. Some questions you can answer from the vendor's website, published studies, and public materials; others should be sent to the vendor for a written response. Record what you find and note the source. Gaps are findings too: if the information is unavailable or the vendor declines to answer, note that.

Pillar 1 — Systemic Signal Assessment

QUESTION	WHAT WE FOUND
P1.Q1 What specific clinical or operational problem does this tool address? What does it do, step by step, from the care team's perspective? *[SEEDS]*	
P1.Q2 What does the vendor say about how it identified this problem — from direct observation of care delivery, staff or patient feedback, or primarily market analysis? Note the source. *[SEEDS]*	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Pillar 2 — Everyone-Centered Design

No pre-meeting research items — this pillar is assessed in Part 2: Evaluation Rubric.

Pillar 3 — Ecosystem Alignment

QUESTION	WHAT WE FOUND
P3.Q1 Does this tool use artificial intelligence (AI), rule-based logic, or a combination? For any AI component: how was the model trained, and how is it updated over time? *[SEEDS]*	
P3.Q2 What changes have been made to the tool in the past 12 months, and what is on the development roadmap for the next 12? *[SEEDS]*	
P3.Q3 What independent evaluations, peer-reviewed studies, or published data support this tool's effectiveness? In what care settings and with what patient populations? *[SEEDS]*	
P3.Q4 How many organizations are currently using this tool? How many are Federally Qualified Health Centers (FQHCs), community health centers (CHCs), or safety-net providers? *[SEEDS]*	
P3.Q5 What outcomes — clinical, operational, and patient experience — have customers documented? What does the data show? *[SEEDS]*	
P3.Q6 Which electronic health record (EHR) systems does this tool integrate with? What is the depth of that integration: read access, write access, bidirectional, or embedded? *[SEEDS] [HSCC]*	
P3.Q7 What device and network requirements does the tool have? Are there bring-your-own-device (BYOD) scenarios customers have navigated? *[SEEDS] [HSCC]*	
P3.Q8 What is the full cost picture beyond licensing: implementation, training, ongoing support, and integration? What does year-three pricing typically look like?	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	WHAT WE FOUND
[SEEDS]	
P3.Q9 What are the contract terms: minimum length, renewal provisions, and price escalation clauses (rules in the contract that allow the vendor to raise prices over time, and what limits exist)? *[SEEDS]*	
P3.Q10 What are the exit provisions? How much notice is required, and what does the exit process require of the health center? *[SEEDS]*	
P3.Q11 What financial assistance, grant-eligible pricing, or phased payment options does the vendor offer to safety-net providers? Has the vendor worked with collective purchasing groups such as the Commonwealth Purchasing Group (CPG)? *[SEEDS]*	
P3.Q12 Is this vendor part of the Commonwealth Purchasing Group (CPG) network? CPG, a Massachusetts League of Community Health Centers company, negotiates pricing for member health centers, so a participating vendor can offer pre-negotiated terms. *[SEEDS]*	
P3.Q13 What does a typical implementation timeline look like for an organization our size? What are the major milestones? *[SEEDS]*	
P3.Q14 What training does the vendor provide at go-live and on an ongoing basis? Who delivers it, and how much time does it require of staff? *[SEEDS]*	
P3.Q15 What data center infrastructure powers the vendor's AI services: their own facility, a co-location provider, or a major cloud platform (such as AWS, Azure, or Google Cloud)? *[SEEDS]*	
P3.Q16 Does the vendor publish an annual sustainability or environmental impact report? If so, where can it be accessed? *[SEEDS]*	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	WHAT WE FOUND
P3.Q17 What percentage of the vendor's data center energy comes from renewable or carbon-free sources? *[SEEDS]*	

Pillar 4 — Data Sovereignty

QUESTION	WHAT WE FOUND
P4.Q1 What security certifications does this tool hold, such as SOC 2 Type II, HITRUST, or FedRAMP (security certifications that show the vendor has passed outside audits for protecting health data)? *[SEEDS] [HSCC]*	
P4.Q2 What patient health information does this tool access, generate, or retain? Where is that data stored, and for how long? *[SEEDS] [HIPAA Security Rule]*	
P4.Q3 What are the vendor's standard Business Associate Agreement (BAA) terms? Does the BAA specifically address AI-related uses of protected health information (PHI) and personal information under Massachusetts state law (201 CMR 17.00)? *[SEEDS] [HIPAA Security Rule]*	

Pillar 5 — Stewardship

No pre-meeting research items — this pillar is assessed in Part 2: Evaluation Rubric.

Part 2: Evaluation Rubric

Use this rubric to judge whether the vendor adequately addressed each area, drawing on your Part 1 findings and what the vendor demonstrates. You may complete it across more than one meeting. For each item, mark Yes, Partially, or Not yet, and record the evidence or source. A STOPGATE item marked Not yet is a binding condition that halts the process until it is resolved.

Pillar 1 — Systemic Signal Assessment

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P1.Q3 The vendor could credibly describe, from the perspective of the people doing the work, how the	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
problem this tool addresses was identified, not just a market rationale. *[SEEDS]*		

Pillar 2 — Everyone-Centered Design

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P2.Q1 The vendor identified the patient populations and care settings this tool was designed and tested for, including populations it is less suited to serve. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q2 The vendor explained how much control the user has over AI-generated outputs and recommendations, including the ability to override, edit, or reject what the tool produces. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q3 The vendor's team — including design, development, and leadership — includes people with lived experience in the communities this tool will serve. The vendor described how patients, frontline staff, or community members are involved in ongoing product decisions, and whether the board or advisory group includes representation from safety-net healthcare or the populations served. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q4 The tool was designed with input from patients, caregivers, or frontline staff in safety-net or community health settings, and the vendor described how that input shaped the current product. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q5 The vendor provided evidence of how this tool performs across patient languages, literacy levels, or other population characteristics. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P2.Q6 The vendor described how the AI model was trained and how it is evaluated for bias across the patient populations this health center serves, including how that performance is monitored after go-live. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 3 — Ecosystem Alignment

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P3.Q18 The vendor described what this tool requires of staff at the point of care and walked through a typical interaction from the staff perspective. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q19 The vendor described the documented limitations of this tool — what it does less reliably or situations where it is less effective. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q20 The vendor provided references from organizations with a patient population and care context similar to ours, including at least one safety-net or FQHC setting. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q21 The vendor disclosed whether any customers have discontinued use of this tool after go-live and explained the reasons. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q22 The vendor described what technical implementation requires from the health center's IT team: estimated staff hours, system access levels, infrastructure, and timeline. *[SEEDS] [HSCC]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q23 The vendor described the fallback when this tool is unavailable or produces an unexpected result, and	☐ Yes ☐ Partially	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
reported how much downtime customers have typically experienced. *[SEEDS] [HSCC]*	☐ Not yet	
P3.Q24 The vendor walked through a complete use case showing whether this tool adds steps to existing care processes, replaces steps, or runs in the background. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q25 The vendor addressed whether this tool creates excessive alerts, interruptions, or added steps in the care workflow, and described how customers have managed alert volume. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q26 The vendor has experience working with Federally Qualified Health Centers (FQHCs) or community health centers (CHCs), and its pricing, implementation approach, and product design reflect an understanding of how FQHCs are funded, including grant funding, sliding fee scales, and payer mix (insurance company reimbursement). *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q27 Where this tool touches billing, coding, or documentation that feeds reimbursement, the vendor disclosed whether it has affected billing or reimbursement outcomes for customers, positive or negative. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q28 The vendor identified additional costs that commonly emerge after go-live, such as expanded licensing, add-on modules, version upgrades, or integration support. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q29 The vendor described its primary funding sources — venture capital, grants, revenue, or a combination — and explained what growth milestones,	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
investor expectations, or revenue targets could affect product development, pricing, or support commitments over the next two to three years. The vendor described how it plans to sustain operations and product support if growth targets are unmet. *[SEEDS]*		
P3.Q30 The vendor described what post-go-live support looks like: the designated contact, typical response time, and what the service-level agreement (SLA) commits to. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q31 The vendor described the most common implementation challenges for organizations similar to ours and how customers have typically resolved them. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q32 The vendor agreed to provide customer references beyond its pre-selected list — for example, drawn from a broader customer list it provides. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 4 — Data Sovereignty

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P4.Q4 The vendor confirmed who owns the data generated through use of this tool: the health center, the vendor, or shared. **STOPGATE** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q5 The vendor confirmed it does not sell, license, or share health center or patient data with any third party for that party's own purposes, beyond the named subprocessors and only to deliver the contracted service. **STOPGATE** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q6 The vendor confirmed the health center must authorize in writing any use of its data to train, refine, or	☐ Yes ☐ Partially	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
improve the vendor's AI models, can opt out, that removing identifying details (de-identification) for AI training requires written authorization, and that the health center shares in the value created through direct compensation or a reduction in cost. **STOPGATE** *[SEEDS]*	☐ Not yet	
P4.Q7 The vendor described what happens to health center and patient data when the contract ends — deletion, return, or archiving — including the timeline and what it requires of the health center. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q8 The vendor identified all third parties — such as subprocessors (other companies the vendor works with to handle or store data), cloud providers, or analytics partners — that have access to patient data, under what terms, and whether those third parties also operate under a BAA or equivalent data protection agreement. *[SEEDS] [HSCC]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q9 The vendor described its breach notification process and what it would disclose to a prospective customer about past security incidents. *[SEEDS] [HSCC]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q10 The vendor demonstrated appropriate safeguards for protecting organizational data, such as encryption, access controls, audit logging, incident response procedures, or vulnerability management. *[SEEDS] [HSCC]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q11 The vendor confirmed the health center retains ownership and control of its data and can see, contest, restrict, or withdraw specific data uses at any time. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 5 — Stewardship

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P5.Q1 The vendor described what governance controls the contract supports — for example, whether the health center can pause, restrict, or decommission the tool at any time without financial penalty. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q2 The vendor reported whether customers have seen impact on quality metric performance, including Healthcare Effectiveness Data and Information Set (HEDIS) measures (standard quality benchmarks used by health insurers and health centers to track patient care outcomes) or Uniform Data System (UDS) reporting (a federal reporting system that FQHCs use to show how well they serve their communities), and identified what measures the health center should plan to monitor after go-live. **[Clinical]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q3 The vendor disclosed whether it has discontinued, significantly reduced support for, or moved away from a product serving healthcare or safety-net populations in the past five years, and described what happened to those customers. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q4 The vendor described how it supports workforce readiness, including the training and upskilling staff need to use the tool well and how that support continues as roles and the tool change over time. *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

Summary of Findings

After completing both parts of this checklist, use this section to record the factual highlights before forwarding findings to the AI Signal Decision Assessment.

SUMMARY AREA	KEY FINDINGS
Strongest vendor evidence	
Gaps or unanswered questions	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

SUMMARY AREA	KEY FINDINGS
Areas marked "Partially" or "Not yet" in Part 2	
STOPGATE binding conditions that remain unresolved	

This summary captures facts, not recommendations. Evaluation of fit happens through the AI Signal Decision Assessment (Section 4), where SEEDS pillars structure the governance decision. Any STOPGATE marked Not yet travels with these findings as a binding condition.

KEY TAKEAWAYS BY ROLE

Leadership

- This checklist gives your AI Governance Hub the facts it needs before any vendor conversation becomes a commitment. Require it before any contract discussion moves forward.
- Treat STOPGATE conditions as non-negotiable. A vendor that cannot meet data ownership, no third-party reselling, or model-training authorization and value return is not ready to proceed.
- When every health center asks these same questions, it sets a consistent standard for how vendors engage with community health centers. Completing this checklist protects your organization and strengthens the field.

AI Governance Hub members

- Complete Part 1 from research and written questions sent to the vendor in advance; keep it current across rounds. Use Part 2 to evaluate the vendor across one or more meetings.
- Record exactly what the vendor says and the evidence for it. Findings, including any STOPGATE conditions, flow directly into the AI Signal Decision Assessment.

IT & Security

- Own the Ecosystem Alignment and Data Sovereignty questions. Flag any answer that is incomplete or inconsistent with your infrastructure.
- Confirm the Data Sovereignty STOPGATE conditions: data ownership, no third-party reselling, and written authorization and value return for model training.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**

This checklist is for when a specific vendor is on the table. If you are still exploring, start in Phase 1: Listen and Learn (Section 3) with the AI Signal Intake Form Template.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

→ Evaluating a specific AI tool?

Complete this checklist before your first vendor meeting. Findings feed the AI Signal Decision Assessment in Phase 2: Assess and Decide (Section 4).

→ Formalizing an existing AI tool?

Use this checklist to capture what the vendor's tool actually does, even after it is already in use. Findings feed the AI Signal Decision Assessment in Phase 2: Assess and Decide (Section 4).

→ Reassessing a changing AI tool?

When a vendor adds features or changes terms, run this checklist again on the changed elements. Compare the new findings against the original to see whether a fresh assessment is needed in Phase 2: Assess and Decide (Section 4).

→ Checking an approved AI tool?

The check itself uses the Stewardship Review Checklist in Phase 4: Steward (Section 6). Return to this checklist if the check reveals the vendor changed something that needs re-capture.

→ Replacing or exiting an AI tool?

For the outgoing tool, exit planning lives in Phase 4: Steward (Section 6). To evaluate a replacement vendor, complete this checklist for the new vendor as part of Phase 2: Assess and Decide (Section 4).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI Signal Decision Assessment

This tool answers: Should we move forward, and how strong is our reasoning?

Phase 2 — Assess and Decide (Section 4)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A five-pillar SEEDS assessment. It earns a reasoning level per pillar, then produces a recommendation.	Work through five pillars in order. Confirm Yes or No, record each earned reasoning level, then recommend.	The AI Governance Hub completes it. Legal counsel reviews Pillar 4. The executive sponsor reviews the recommendation.

Before You Begin

A note on time. Running this assessment well takes more than one meeting. For most signals, the AI Governance Hub completes Pillars 1 and 2 in one session, gathers what it needs for the remaining pillars, and returns over two or three cycles. The recommendation is complete when the AI Governance Hub has honest answers across all five pillars.

For who sits on the Hub and each role, see the AI Governance Hub Getting Started Kit (Section 2).

> **LEGAL NOTE**
>
> This assessment is a governance resource, and legal counsel provides legal advice. Each health center is responsible, with its own legal counsel, for reviewing AI vendor contracts, confirming HIPAA and state law compliance, and ensuring data practices are lawful. These questions guide the AI Governance Hub's reasoning. Legal counsel answers them for a specific situation.

Bracketed source keys identify the framework each question draws on; full citations appear in Appendix B: Implementation Approach Reference. Every question in this assessment draws on SEEDS of Innovation™ [SEEDS].

Signal Summary

Transfer from the AI Signal Intake Form Template: description, source, and date. Then record the entry scenario that fits best; the six entry scenarios are in Section 1.2.

Signal description	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Source of signal	
Date brought to AI Governance Hub	
Entry scenario	
Completed by	
Date completed	

How Scoring Works

Each pillar is assessed across five reasoning-maturity levels — Reflexive, Protective, Normative, Integrative, Regenerative — that measure how deeply the AI Governance Hub listened, not only whether a box is checked. Confirm Yes or No for each question and record the evidence. Levels are cumulative: your earned level is the highest level where every question at that level, and every level below it, is answered Yes. Not every question will be answerable; an honest No names where the Hub needs to listen more deeply.

Worked Example

Signal description	The electronic health record (EHR) vendor activated an AI note summarization feature in the latest update. Staff learned of it after it appeared. Two providers flagged it during the morning huddle.
Source of signal	Staff concern (Track 2, AI Signal Intake Form Template)
Entry scenario	AI arrived without a formal review
Completed by	Maria Santos, AI Governance Hub Lead

Pillar 1: Systemic Signal Assessment

What problem does this actually solve, and is that the right problem to address?

QUESTION	YES / NO	EVIDENCE
Level 1: Reflexive — Catch the signal		
P1.Q1 Have we identified the specific problem or opportunity that brought this signal to the AI Governance Hub, and can we describe the ongoing pattern? *[SEEDS]*	☐ Yes ☐ No	
Level 2: Protective — Name the risks		
P1.Q2 Have we assessed whether this signal creates liability, compliance, or security exposure? *[SEEDS]*	☐ Yes ☐ No	
Level 3: Normative — Follow the process		

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / NO	EVIDENCE
P1.Q3 Have we heard directly from the people closest to this problem, and has what they said shaped how we understand it? *[SEEDS]*	☐ Yes ☐ No	
P1.Q4 Have we mapped the underlying pattern that produces this outcome, beyond the visible symptom? *[SEEDS]*	☐ Yes ☐ No	
Level 4: Integrative — See the full picture		
P1.Q5 Have we confirmed that AI is the right solution, rather than a policy change, workflow redesign, or staffing adjustment? *[SEEDS]*	☐ Yes ☐ No	
P1.Q6 Have we checked whether something like this has been tried before, here or elsewhere in the network, and documented what happened? *[SEEDS]*	☐ Yes ☐ No	
Level 5: Regenerative — Build for the future		
P1.Q7 Have we defined what success looks like in 12 months, in concrete and measurable terms? *[SEEDS]*	☐ Yes ☐ No	

Earned reasoning level reminder: *Your earned level is the highest level where all questions at that level and every level below it are answered Yes.*

Pillar earned reasoning level: Level ___

Pillar 2: Everyone-Centered Design

Who was in the room when this was designed? Who benefits, and who carries the risk?

QUESTION	YES / NO	EVIDENCE
Level 1: Reflexive — Catch the signal		
P2.Q1 Can we name the people this tool will affect at this health center — whether patients, staff, or both? *[SEEDS]*	☐ Yes ☐ No	
Level 2: Protective — Name the risks		
P2.Q2 Have we identified who carries more burden or risk from this tool, and why — including people who face the most barriers (language, access, role, digital literacy)? *[SEEDS]*	☐ Yes ☐ No	
Level 3: Normative — Follow the process		

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / NO	EVIDENCE
P2.Q3 Has this tool been tested or validated for the people it will serve or support at this health center — across the conditions, languages, roles, and settings that apply? *[SEEDS]*	☐ Yes ☐ No	
P2.Q4 Have the people who would use or be affected by this tool been included in evaluating it, and has their input been documented? *[SEEDS]*	☐ Yes ☐ No	
Level 4: Integrative — See the full picture		
P2.Q5 Have we brought existing feedback from the people affected — from the Patient AI Feedback Form Template, staff concerns through the AI Signal Intake Form Template, or other sources — into this assessment as direct evidence? *[SEEDS]*	☐ Yes ☐ No	
P2.Q6 Have we assessed whether this tool could reduce access, create confusion, or shift burden onto the people who already face the most barriers? *[SEEDS]*	☐ Yes ☐ No	
Level 5: Regenerative — Build for the future		
P2.Q7 Have we designed a plan to monitor this tool's impact across different groups after go live, with named owners and defined review points? *[SEEDS]*	☐ Yes ☐ No	

Earned reasoning level reminder: *Your earned level is the highest level where all questions at that level and every level below it are answered Yes.*

Pillar earned reasoning level: Level ___

Pillar 3: Ecosystem Alignment

Can the real system hold this tool without adding burden to people who are already at capacity? Confirm the fit — operational and financial — before you commit.

QUESTION	YES / NO	EVIDENCE
Level 1: Reflexive — Catch the signal		
P3.Q1 Have we assessed how this tool changes the daily work for the people closest to it — does it remove steps, add steps, or shift how work gets done? *[SEEDS]*	☐ Yes ☐ No	
Level 2: Protective — Name the risks		

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / NO	EVIDENCE
P3.Q2 Have we assessed what implementation requires in staff time, training time, IT resources, and leadership attention — and whether the organization has genuine capacity right now? *[SEEDS]*	☐ Yes ☐ No	
Level 3: Normative — Follow the process		
P3.Q3 Have we assessed what happens when this tool fails or produces an incorrect output — and confirmed the organization can continue operating without it? *[SEEDS]*	☐ Yes ☐ No	
P3.Q4 If this tool requires integration with existing systems, have we confirmed the technical fit is feasible and resourced? *[SEEDS]*	☐ Yes ☐ No	
Level 4: Integrative — See the full picture		
P3.Q5 Have we assessed the financial fit — whether the health center's margins can absorb the cost, and whether reimbursement and quality incentives support its use? *[SEEDS]*	☐ Yes ☐ No	
P3.Q6 Have we assessed whether this tool adds cognitive burden — more mental load without removing equivalent load — and whether the alignment rests on a workaround rather than a genuine fit? *[SEEDS]*	☐ Yes ☐ No	
Level 5: Regenerative — Build for the future		
P3.Q7 Have we assessed how this tool affects the broader ecosystem — including existing vendor contracts, payer requirements, and regulatory alignment — and confirmed compatibility? *[SEEDS]*	☐ Yes ☐ No	

Earned reasoning level reminder: *Your earned level is the highest level where all questions at that level and every level below it are answered Yes.*

Pillar earned reasoning level: Level ___

> **IN PLAIN TERMS**
>
> Cognitive burden means the tool adds mental load without removing it: more steps, more alerts, more decisions layered onto an already full workflow. A tool that rests on a workaround, where a

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> staff member carries an extra step because they want it to succeed, is a tool whose alignment is unconfirmed.

Pillar 4: Data Sovereignty

Who owns the data, and what rights do the health center and the people it serves have over how it is used?

QUESTION	YES / NO	EVIDENCE
Level 1: Reflexive — Catch the signal		
P4.Q1 Do we understand what data this tool uses, what data it generates, and where that data goes across its full lifecycle? *[SEEDS]*	☐ Yes ☐ No	
Level 2: Protective — Name the risks		
P4.Q2 Have we confirmed who owns the data this tool generates — the health center, the vendor, or both — and is this explicit in the contract? *[SEEDS]*	☐ Yes ☐ No	
P4.Q3 Do we know what secondary uses of data the vendor retains rights to — analytics, benchmarking, model training — and has the AI Governance Hub assessed whether these are acceptable? *[SEEDS]*	☐ Yes ☐ No	
Level 3: Normative — Follow the process		
P4.Q4 Has the health center's legal counsel reviewed the vendor contract and any Business Associate Agreement (BAA) covering this tool, and confirmed compliance with HIPAA, HITECH (the federal law that strengthened HIPAA's rules on electronic health records and data breaches), and applicable state privacy requirements? *[SEEDS]*	☐ Yes ☐ No	
Level 4: Integrative — See the full picture		
P4.Q5 Have we assessed what happens to the data if the vendor relationship ends — can the health center retrieve it, and can it be deleted from the vendor's systems? *[SEEDS]*	☐ Yes ☐ No	
Level 5: Regenerative — Build for the future		

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / NO	EVIDENCE
P4.Q6 Have we assessed whether the data governance for this tool strengthens the health center's long-term data stewardship — including access controls, accountability, and the trust of the people whose data it holds? *[SEEDS]*	☐ Yes ☐ No	

Earned reasoning level reminder: *Your earned level is the highest level where all questions at that level and every level below it are answered Yes.*

Pillar earned reasoning level: Level ___

> **LEGAL NOTE**
> Each community health center remains solely responsible for the accuracy of claims it submits, the security of patient health information, and the integrity of its data governance practices, regardless of whether AI was involved. Legal review of contracts is required before any tool is approved.

Pillar 5: Stewardship

Who is accountable for this tool after it is live, and every year after, and how will you know whether it is still doing what it was intended to do?

QUESTION	YES / NO	EVIDENCE
Level 1: Reflexive — Catch the signal		
P5.Q1 Have we named the person inside the health center who is accountable for this tool after it is live (the person responsible for confirming it keeps working as intended and who will act if something goes wrong)? *[SEEDS]*	☐ Yes ☐ No	
Level 2: Protective — Name the risks		
P5.Q2 Have we defined what would cause the health center to pause or stop this tool — specific conditions, thresholds, or events — and does the accountable person have the authority to act? *[SEEDS]*	☐ Yes ☐ No	
P5.Q3 Have we assessed how the AI Governance Hub will identify emerging harms — including harms that may only become visible after the tool has been in use, such	☐ Yes ☐ No	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / NO	EVIDENCE
as changes in how staff work, shifts in who benefits, or effects the vendor did not predict? *[SEEDS]*		
Level 3: Normative — Follow the process		
P5.Q4 Has the AI Governance Hub set a formal review schedule for this tool — including what the review covers, who participates, and what would trigger a review outside the regular cadence? *[SEEDS]*	☐ Yes ☐ No	
Level 4: Integrative — See the full picture		
P5.Q5 Have we defined what success looks like — using specific measures across tool performance, the experience of the people using it, and the experience of the people it affects — and confirmed those measures with the people closest to the work? *[SEEDS]*	☐ Yes ☐ No	
P5.Q6 Have we assessed how feedback from the people using or affected by this tool will be collected, reviewed, and acted on — so that concerns become governance signals rather than unresolved complaints? *[SEEDS]*	☐ Yes ☐ No	
Level 5: Regenerative — Build for the future		
P5.Q7 Have we assessed how this decision strengthens the health center's governance capacity over time — so that each AI decision makes the next one easier, better informed, and less fragile? *[SEEDS]*	☐ Yes ☐ No	

Earned reasoning level reminder: *Your earned level is the highest level where all questions at that level and every level below it are answered Yes.*

Pillar earned reasoning level: Level ___

Pillar Reasoning Summary

Record the earned reasoning level from each pillar. Levels are cumulative: your earned level is the highest level where all questions at that level and every level below it are answered Yes. Completed by the AI Governance Hub Lead or designated working group lead after all five pillars are assessed.

Pillar 1: Systemic Signal Assessment	Level ___
Pillar 2: Everyone-Centered Design	Level ___

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Pillar 3: Ecosystem Alignment	Level ___
Pillar 4: Data Sovereignty	Level ___
Pillar 5: Stewardship	Level ___

IMPORTANT

The recommendation should reflect the full profile. A gap named is stronger than a gap hidden. The lowest pillar matters most: a tool at Level 4 on four pillars and Level 2 on one has a gap the AI Governance Hub must address.

AI Governance Hub Recommendation

Record the recommendation based on the five-pillar assessment and the reasoning summary. This recommendation feeds directly into the Design Plan Template (Section 5: Phase 3 — Plan and Implement), where the decision, conditions, and plan are documented.

☐ Approve	The AI Governance Hub answered all five pillars with sufficient confidence. The reasoning profile supports moving forward.
☐ Approve with conditions	The AI Governance Hub identified specific gaps that must be resolved as the tool moves forward. Conditions are documented in the Design Plan Template.
☐ Hold	The AI Governance Hub identified gaps that must be resolved before a recommendation can be made. The signal remains open.
☐ Do not proceed	The assessment found a concern the AI Governance Hub cannot resolve. The reasoning is documented. This signal may return when conditions change. When the concern is with the specific market option rather than the need itself, the need can re-enter as a new signal, including a build path, and earn a fresh assessment here.

Reasoning summary	Summarize the AI Governance Hub's reasoning in plain language. Reference which pillars were strong, which had gaps, and how the reasoning profile informed the recommendation. Honest and brief is the standard.
Unresolved gaps	List any gaps the AI Governance Hub named but could not resolve during this assessment. These carry forward into the Design Plan Template.

IMPORTANT

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> Buy first, then build. This tool assumes the health center looks first at what it can buy: whether a product already exists in the market that meets the need. Use the Vendor Evaluation Checklist to gather the facts and confirm fit. When nothing available to buy meets the need, the health center looks at building the capability itself. When the choice is not yet clear, record what is still missing and who will find out.

Next step. Document this recommendation in the Design Plan Template (Section 5: Phase 3 — Plan and Implement). The Design Plan Template records the decision, names conditions and owners, builds the implementation plan, and maintains the living record for ongoing review.

KEY TAKEAWAYS BY ROLE

Leadership	AI Governance Hub members	IT & Security
• A completed AI Signal Decision Assessment comes before any AI tool recommendation is finalized. Your review confirms the AI Governance Hub applied structured reasoning. • The reasoning profile across five pillars matters as much as the recommendation itself. Review the profile, especially the lowest pillar.	• This is a working document. Complete it across multiple sessions. Honest gaps are more valuable than checked boxes. • Levels are cumulative. A No at any level means the AI Governance Hub resolves it before claiming a higher level.	• Pillar 3 (Ecosystem Alignment) and Pillar 4 (Data Sovereignty) require your input on technical fit, data flows, and vendor contract terms. • Coordinate with the AI Governance Hub Lead to confirm integration feasibility and security posture before the recommendation is recorded.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

Complete this assessment to reach a recommendation, then document it in the Design Plan Template (Section 5: Phase 3 — Plan and Implement). If approved, complete the Pre-Deployment Impact Checklist before go-live.

→ Evaluating a specific AI tool?

Complete the Vendor Evaluation Checklist alongside this assessment, then document the decision in the Design Plan Template (Section 5: Phase 3 — Plan and Implement).

→ Formalizing an existing AI tool?

Assess the tool already in use. The recommendation determines whether it stays, changes, or is removed. Document it in the Design Plan Template.

→ Reassessing a changing AI tool?

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Reassess under the new conditions to confirm whether the original approval still holds. Document the update in the Design Plan Template.

→ Checking an approved AI tool?

Start with the Stewardship Review Checklist (Section 6: Phase 4 — Steward). If the review surfaces a concern that needs a fresh governance assessment, return here.

→ Replacing or exiting an AI tool?

Manage the exit through the Stewardship Review Checklist (Section 6: Phase 4 — Steward). A replacement tool enters as a new signal and runs through this assessment.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Pre-Deployment Impact Checklist

This tool answers: Is this AI tool safe to deploy?

Phase 2 — Assess and Decide (Section 4)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A readiness checklist organized by the five SEEDS pillars, for clinical or administrative AI tools, bought or built.	Answer each question Yes, Partially, or Not yet with evidence. A STOPGATE item at Partially or Not yet holds launch.	The Hub Lead assigns each pillar to the right expertise; the signed, evidenced checklist becomes the formal go-live record.

Before You Begin

The AI Governance Hub Lead coordinates this checklist and assigns each pillar to the member with the most relevant expertise. Before starting, gather the results of the AI Signal Decision Assessment for this tool, any vendor-supplied documentation (contract, Business Associate Agreement, technical specifications), training completion records, and the rollout and rollback plans.

Think of this checklist as the safety check for the decision: it confirms documented evidence that the tool is safe to deploy. The companion check at the moment of launch, confirming that people, patients, and launch-day safeguards are ready, is the AI Implementation Readiness Checklist (Phase 3 — Plan and Implement, Section 5). A Yes answer here always rests on documented evidence, so the record holds up under review.

IMPORTANT

EVIDENCE OVER ASSERTION: A Yes answer requires documented evidence — a signed record, a completed test log, or a verified result. A Partially answer means the work is underway and documentation is incomplete: record what exists and what remains. Vendor assurances alone are insufficient. When documentation is still pending, the answer is Not yet.

Reading the Scope Tags

Each question ends with one or more tags.

Acquisition — which path the tool came from:

- **[Buy]** applies when the health center purchased the AI tool from a vendor.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- **[Build]** applies when the health center developed the tool internally.
- **[Buy or Build]** applies in every case.

Context — where the tool operates:

- **[Clinical]** applies when the AI tool supports direct patient care or clinical decisions.
- **[Admin]** applies when the AI tool supports administrative or operational functions.
- No context tag means the question applies in both settings.

Binding condition:

- **STOPGATE** halts the go-live process until it is resolved, because going forward would create harm or exposure that cannot be undone after the fact.

Bracketed source keys identify the framework each question draws on; full citations appear in Appendix B: Implementation Approach Reference.

Systemic Signal Assessment

These questions read what this tool is and the risk it carries, before the checklist goes further. The tool's purpose is first recorded in the AI Signal Intake Form Template (Section 3) and the AI Signal Decision Assessment; this pillar confirms that record.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P1.Q1 Has the AI Governance Hub documented what this tool does, who will use it, and what decisions it supports? **[Buy or Build]** *[HSCC Phase 0] [NIST MAP-2]*	☐ Yes ☐ Partially ☐ Not yet	
P1.Q2 Has the AI Governance Hub assigned a risk level (low, medium, high, or critical) based on the decisions this tool influences, the sensitivity of the data involved, and the population it affects? **[Buy or Build]** *[HSCC Phase 0] [NIST MAP-1]*	☐ Yes ☐ Partially ☐ Not yet	
P1.Q3 Has the company providing the tool supplied written documentation of the tool's intended use, its known limitations, and the populations or care settings where it has been tested? **[Buy]** *[HSCC Phase 0] [NIST MAP-3]*	☐ Yes ☐ Partially ☐ Not yet	
P1.Q4 Has the internal team documented the data sources used to develop this tool, the development	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
approach, known limitations, and intended scope? **[Build]** *[HSCC Phase 0] [NIST MAP-3]*		

Everyone-Centered Design

These questions confirm the tool works for the people it serves, especially those who face the most barriers to care, and for the staff who use it.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P2.Q1 Has the tool been evaluated for differential performance across the populations this health center serves, with documented results? **[Buy or Build]** *[NIST MEASURE-2] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q2 Has training been documented as complete for all staff who will use this tool, in the languages and formats they need, with a documented channel for raising concerns during the transition? **[Buy or Build]** *[HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q3 Do patients receive clear, plain-language information about how AI is used in their care before their first encounter with this tool, in the primary languages of the patients this tool will serve? **[Buy or Build]** **[Clinical]** **STOPGATE** *[HSCC Phase 3] [ONC SAFER]*	☐ Yes ☐ Partially ☐ Not yet	

Ecosystem Alignment

These questions confirm the tool holds across the whole system it joins — clinicians, staff, existing systems, regulators, and the capacity to run it. Vendor-supplied evidence is captured in the Vendor Evaluation Checklist and confirmed here.

The human-oversight floor every AI tool operates under is set in the AI Policy Checklist.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P3.Q1 Has the tool been tested in a care or operational setting comparable to this health center, with documented results showing it performs as described? **[Buy or Build]** *[HSCC Phase 3] [NIST MEASURE-2]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q2 Has this tool received the regulatory authorization required for its intended clinical use, including Food and Drug Administration (FDA) clearance where applicable? **[Buy or Build]** **[Clinical]** **STOPGATE** *[ECRI TSS] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q3 Does the tool clearly label AI-generated content and provide an explanation of how it reached its output, so the person using it can make an informed decision? **[Buy or Build]** *[ONC SAFER]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q4 Can clinicians freely override or set aside AI-generated recommendations, keeping their standard process and documentation intact? Are clinical alerts calibrated to a volume and relevance level that avoids alert fatigue? **[Buy or Build]** **[Clinical]** *[HAIRA] [ONC SAFER]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q5 Can every person who will use this tool answer three questions before go-live: (1) what the tool does, (2) what it does well, and (3) when to override or set aside its output? **[Buy or Build]** *[HAIRA] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q6 Do all staff who will use this tool have the devices, system access, and credentials they need before go-live? **[Buy or Build]** *[ONC SAFER] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q7 Has the tool been tested for integration with the health center's existing systems (electronic health record, scheduling, billing, or other platforms), with documented results? **[Buy or Build]**	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
[HSCC Phase 3] [ONC SAFER]		
P3.Q8 Are staffing levels adequate to manage the go-live transition while maintaining current care capacity? **[Buy or Build]** *[ONC SAFER] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q9 Is a phased rollout plan documented, with a defined scope for initial launch and clear criteria for expanding to full use? **[Buy or Build]** *[HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	

Data Sovereignty

These questions confirm the data the tool touches stays protected and governed, with clear agreements over who handles it.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P4.Q1 Is the signed Business Associate Agreement (BAA) on file before go-live, confirming it covers all data this tool will access, process, or store and names every subcontractor identified during vendor evaluation? **[Buy]** **STOPGATE** *[HSCC Phase 2] [NIST GOVERN-6]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q2 Has the internal team completed a security review of the tool and tested it in a staging environment, with documented baseline performance, before moving to production? **[Build]** *[HSCC Phase 3] [NIST GOVERN-1]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q3 Has IT confirmed, in the health center's own environment, that HIPAA compliance, data encryption, and activity logging are configured and working before go-live? **[Buy or Build]** **STOPGATE** *[HSCC Phase 3] [ONC SAFER]*	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P4.Q4 Do patients receive clear, plain-language information about how their data from this tool is used and where it goes, including any use beyond their direct care, before their first encounter, in the primary languages of the patients this tool will serve? **[Buy or Build]** **[Clinical]** *[SEEDS] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	

Stewardship

These questions confirm the health center can watch the tool after go-live and keep the power to stop it. The rollback plan is a binding condition; the rest are required with a documented resolution plan.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P5.Q1 Can the decisions this tool makes or supports be audited, with a clear record of what the AI recommended and what action was taken? **[Buy or Build]** **[Admin]** *[NIST GOVERN-1] [HSCC Phase 3]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q2 Is a rollback plan documented and ready, so the health center can return to the previous process if the tool fails or causes harm? **[Buy or Build]** **STOPGATE** *[HSCC Phase 3] [ECRI TSS]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q3 Has the AI Governance Hub defined what will be measured after go-live, how often reviews will occur, and who is responsible for reviewing the results? **[Buy or Build]** *[NIST MEASURE-3] [ECRI TSS]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q4 Is an incident response plan in place that defines how the health center will respond if the tool produces an error, a safety concern, or an unexpected result? **[Buy or Build]** *[HSCC Phase 3] [ECRI TSS]*	☐ Yes ☐ Partially ☐ Not yet	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Decision Summary

Record how many questions in each pillar are answered Yes, Partially, or Not yet, and note where evidence is still missing.

PILLAR	YES	PARTIALLY	NOT YET	EVIDENCE GAPS
Systemic Signal Assessment (4)	___	___	___	
Everyone-Centered Design (3)	___	___	___	
Ecosystem Alignment (9)	___	___	___	
Data Sovereignty (4)	___	___	___	
Stewardship (4)	___	___	___	
STOPGATE binding conditions (5)	___	___	___	☐ All cleared ☐ Unresolved

Binding conditions (STOPGATE), one per sign-off role: patient disclosure (P2.Q3, Patient Access/Experience Lead); regulatory authorization (P3.Q2, Clinical Leadership); Business Associate Agreement (P4.Q1, Legal / Compliance); HIPAA, encryption, and logging (P4.Q3, IT and Security Lead); rollback plan (P5.Q2, AI Governance Hub Lead). A Partially or Not yet on any binding condition holds go-live.

Go-Live Decision (select one):

☐ Go Live — all questions answered Yes with documented evidence; all STOPGATE binding conditions cleared.

☐ Conditional Go-Live — all STOPGATE binding conditions cleared; Partially answers elsewhere have documented resolution plans with owners and deadlines.

☐ Limited Pilot Only — one or more STOPGATE binding conditions marked Partially; scope restricted to a defined pilot pending resolution.

☐ Hold — one or more STOPGATE binding conditions marked Not yet with an unclear resolution path; return to the checklist when resolved.

Sign-Off

All designated signers must sign before a go-live decision is recorded and acted on.

ROLE	PRINT NAME	SIGNATURE	DATE
AI Governance Hub Lead			
Clinical Leadership			
IT and Security Lead			
Legal / Compliance			

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ROLE	PRINT NAME	SIGNATURE	DATE
Patient Access/Experience Lead			

KEY TAKEAWAYS BY ROLE

Leadership

- This checklist is the health center's formal record of AI readiness before go-live; the signed record protects the organization, the board, and patients.
- A Hold or Limited Pilot decision is a governance strength: the process caught a gap before it reached patients or staff.

AI Governance Hub members

- Assign each pillar to the member with the right expertise; the checklist is a coordinated team effort.
- Every STOPGATE binding condition needs a verifiable document, record, or test result before it clears.

IT & Security

- Ecosystem Alignment and Data Sovereignty are yours: confirm integration, security compliance, and data handling before go-live.
- The rollback plan in Stewardship needs your confirmation that the health center can return to the previous process within a defined timeframe.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?

This checklist applies once a specific tool is chosen. Start with the AI Signal Intake Form Template (Phase 1 — Listen and Learn, Section 3); return here when you have a named tool ready for go-live review.

→ Evaluating a specific AI tool?

Complete the Vendor Evaluation Checklist and AI Signal Decision Assessment first (Phase 2 — Assess and Decide, Section 4); bring their evidence here to confirm go-live readiness.

→ Formalizing an existing AI tool?

Run this checklist on the tool already in use. The pillar results show whether it continues as is, continues with conditions, or moves toward removal.

→ Reassessing a changing AI tool?

The tool gained features or expanded use. Re-run this checklist on the changed scope and compare to the original to see where conditions shifted.

→ Checking an approved AI tool?

For a live tool, use the Stewardship Review Checklist (Phase 4 — Steward, Section 6). If the review finds changes large enough to require re-assessment, return here.

→ Replacing or exiting an AI tool?

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Run this checklist on the incoming tool. For exit planning on the outgoing tool, see the Stewardship Review Checklist (Phase 4 — Steward, Section 6).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 5: Phase 3 — Plan and Implement

GOAL OF THIS SECTION

Turn an approved AI decision into a governed plan, a ready organization, and a launch your staff and patients can trust.

WHAT'S COVERED IN THIS SECTION

5.1 From Decision to Designing a Plan
Carry an approved decision into a concrete plan

5.2 Components of a Good Plan
What every plan needs before moving forward

5.3 Workforce Lens and Role Changes
Support staff through the changes AI brings

5.4 Communication with Staff and Patients
Tell staff and patients clearly, before launch

5.5 Implementation Readiness
Confirm readiness at the go-live gate

TOOLS YOU'LL USE

Design Plan Template
How do we design a plan for implementation that works for our patients, our staff, and our health center?

AI Implementation Readiness Checklist
Is our organization ready to launch this?

5.1 From Decision to Designing a Plan

By the time your team reaches this phase, you have already done the hard work of reviewing an AI tool or idea — a signal, meaning any indication, from a person, a system, a vendor, a regulator, or a peer organization, that something AI-related needs attention. Now the question is simpler and larger: are we moving forward, and what is our plan?

The AI Signal Decision Assessment in Phase 2 — Assess and Decide (Section 4) produced one of four outcomes: Approve, Approve with conditions, Hold, or Do not proceed. A decision is complete when it is documented: the outcome named, the reasoning recorded, and the next step assigned.

For an approved signal, the Design Plan Template carries that decision into a concrete plan. It answers one question: how do we design a plan for implementation that works for our patients, our staff, and our health center? It anchors the plan to the decision — the signal, the decision date, the conditions placed on the approval, and any gaps the AI Governance Hub, your team that reviews AI decisions, named but could not resolve. Those gaps carry forward as work this plan must address.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Every documented plan becomes part of your governance record. When staff leave, the record keeps that knowledge in place, visible in the record rather than held in one person's memory. When the same question returns, your Hub can look back instead of starting over. Shared through Massachusetts League of Community Health Centers (MLCHC) channels, one health center's plan becomes something the whole network can learn from.

> **IMPORTANT**
>
> In the Design Plan Template, this information lives in Part 1: Decision Reference — the signal name, decision date, decision type, conditions or limits on the approval, and the unresolved gaps carried forward.

5.2 Components of a Good Plan

A good plan is specific. The Design Plan Template structures every plan around the same parts, so anyone reading it can see what problem is being addressed, who is affected, what will change, who is responsible, what guardrails hold, how success will be measured, and what would cause the health center to stop.

Before a plan moves forward, the AI Governance Hub should be able to answer five questions: What are we trying to improve? Who is affected if this works, or if it does not? What are the guardrails that cannot be crossed? What evidence will tell us whether to expand, adjust, or stop? Who owns this after launch? If any question stays open, the plan is not ready.

The plan measures the tool against the same five SEEDS pillars the AI Signal Decision Assessment used to reach the decision: Systemic Signal Assessment, Everyone-Centered Design, Ecosystem Alignment, Data Sovereignty, and Stewardship. The Design Plan's Evidence Plan sets at least one measure for each pillar that applies — for example, note-completion time with no rise in after-hours work (Ecosystem Alignment), or steady visit access across the languages and populations you serve (Everyone-Centered Design). The Stewardship Review Checklist in Phase 4 — Steward (Section 6) tracks these measures once the tool is live.

When a plan includes a pilot, a limited test before full rollout, the same governance applies. The Design Plan's Pilot Charter sets scope, timeline, measures and targets, early-stop triggers, and a midpoint review. A plan that clears the five questions, sets its guardrails including a cybersecurity review, and defines exit criteria is ready for the readiness check in Section 5.5.

> **IMPORTANT**

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

> These parts live in Part 2: The Plan and the Pilot Charter — ownership, timeline, training, the evidence plan, guardrails, exit criteria, and pilot design. Part 3: Living Record tracks how the plan performs after launch.

5.3 Workforce Lens and Role Changes

Every AI tool changes how at least some people do their work. The change may be small, one fewer step, or larger: a role that needs new skills, or a task where a person reviews what AI produced instead of creating it from scratch.

The Design Plan Template names these changes before go-live. Five questions shape the workforce part of the plan: What processes does this tool touch, and for whom? What does a person in that role need to know to use it well? Who provides training, and how long does it take? How does the tool change what a person is responsible for, and is that clear to them? What happens when the tool produces an error, and who catches it? That last question applies to every tool whose output your health center acts on: who reviews it, what they look for, and how quickly, all named before launch.

Workforce investment travels with tool investment. When a tool is funded, the reskilling it requires is planned and funded alongside it. Where training resources are limited, an AI champion can make the difference between adoption and frustration.

> **IN PLAIN TERMS**
>
> An AI champion is a staff member who learns the tool early and helps colleagues get comfortable with it. This person is a go-to resource for questions and a bridge between the AI Governance Hub and the staff using the tool every day.

By the time a tool completes governance, some staff are likely already using AI on their own: tools they find helpful, apps on personal devices, or AI features switched on inside software you already use. This is normal; people find tools that help them work. Understanding what staff already use shows the Hub where current systems fall short and which tools belong in review. The AI Signal Intake Form Template in Phase 1 — Listen and Learn (Section 3) creates that path: staff submit the tool, the Hub reviews it, and if it meets your criteria it joins the managed inventory. These are tools already in use without review, not violations; bringing them into governance is how the system catches up.

5.4 Communication with Staff and Patients

Clear communication before a tool goes live reduces confusion, builds trust, and surfaces concerns early, while they can still be addressed. The Design Plan Template's Communication Plan names who needs to know, what they need to understand, when, and how.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Staff should know, before go-live: what the tool does, in plain terms; how their day-to-day work changes; what training is available and when; who to contact with questions; and that there is a safe way to raise a problem. That last point matters most in early rollouts. Staff who feel they cannot surface a concern will wait until it grows. A named listening channel, used consistently, turns small problems into information the AI Governance Hub can act on. Concerns may reflect the energy and resource cost of AI, effects on human judgment, professional values, or personal values; each is worth understanding, not only resolving. Workforce policies that intersect with human resources (HR) practices, union agreements, or performance are developed with your HR leadership and your legal and compliance advisors. This Playbook frames the conversation; your organization sets the policy.

Patients affected by an AI tool are told before it is live, especially for tools that touch clinical communication, documentation, or care decisions. Patient communication is honest, in plain language, and ahead of the change: what the tool does, who can access the information it generates, and how to raise a question. Patients are told, in plain language, when AI is used in their care, and every audience is reached in the languages and formats they can understand.

5.5 Implementation Readiness

Before an approved AI tool goes live, your health center confirms it is ready. The AI Implementation Readiness Checklist answers one question: is our organization ready to launch this? The Design Plan built and gated the plan; this checklist verifies what that plan set in motion, so each item confirms readiness rather than re-collecting what the plan already records.

The checklist is organized by four of the five SEEDS pillars. The systemic problem and scope were settled upstream, in the AI Signal Decision Assessment and the Design Plan, so readiness begins with the people the launch touches. Everyone-Centered Design confirms patients can decline the tool with their care unchanged and the patient advisory board is briefed. Ecosystem Alignment confirms staff know the tool is coming, have practiced it in a non-production setting, and know what to do if it fails. Data Sovereignty confirms patients are asked before any recording and that data use is explained in plain language. Stewardship confirms a named change agent can halt the go-live and that escalation paths are set. Each item is marked Yes, Partially, or Not yet.

Four items are STOPGATEs: binding conditions that must read Yes before launch, because going ahead would create harm that cannot be undone. Four are owned by named signers — the patient opt-out, the live human review of anything the tool creates before it enters the record or billing, the change agent's authority to halt, and patient consent before recording. Where any item reads Not yet, the Hub records it with an owner and a date, then chooses Go live, Conditional go-live, or Hold. Four roles sign the decision: Compliance or Legal, Clinical lead, Operations or Implementation lead, and Executive sponsor. When every binding condition reads Yes, the launch proceeds.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

IMPORTANT

A Not yet you can see is one you can act on. If your team cannot answer a readiness question at all, especially a staff-readiness item, that is the clearest sign more preparation is needed before the tool starts.

CASE EXAMPLE: BAYSIDE PLANS AND LAUNCHES THE DOCUMENTATION TOOL

This example is hypothetical and depicts a composite community health center. It continues the Bayside story from Section 4.

The story so far. Bayside Community Health Center's **AI Governance Hub** gathered signals about documentation burden (Section 3) and evaluated a vendor's ambient documentation tool, a tool that listens during the visit and drafts the clinical note (Section 4). The AI Governance Hub approved the tool with conditions and limited it to a small pilot.

The situation. Bayside has said yes, with conditions. The approval from Section 4 carried conditions and one open gap: confirm the vendor's data terms, deliver Portuguese patient-disclosure materials, and settle who reviews each AI-drafted note before it enters the record.

An approval is a decision. This section's two tools turn it into a launch.

Step 1 — Design Plan Template *(How do we design a plan for implementation that works for our patients, our staff, and our health center?)*

The AI Governance Hub members responsible for carrying out the decision start with Part 1, Decision Reference: they copy in what was approved, when, the decision type (New AI tool), the conditions, and the unresolved gap, so nothing gets lost between deciding and doing.

Then they build Part 2, The Plan, answering one plain question at a time. Who owns this? A named person for each perspective: executive sponsor, operations leadership, IT or health informatics, compliance or legal, human resources, and communications. What changes for staff? Clinicians will review a drafted note instead of writing one from scratch, and the plan names who holds that review. What do we tell people? A plain-language patient notice in the languages served, and a named route for staff feedback.

How will we know it works? The Evidence Plan sets one measure for each SEEDS pillar that applies: visit access stays steady across languages, note time goes down without after-hours work going up, every privacy agreement stays in scope, every scheduled review happens on time. Where are the limits? Approved scope is adult primary care at one site, with eight exit criteria that must all read Yes before full rollout.

And because a pilot is planned, the Pilot Charter sets it out: 90 days, one site, with early-stop triggers spelled out, starting with any patient safety event.

Step 2 — AI Implementation Readiness Checklist *(Is our organization ready to launch this?)*

This checklist confirms the last step before launch: that people and patients are ready, and that the launch-day safeguards are live. The AI Governance Hub works through it, marking Yes, Partially, or Not yet.

Four items are STOPGATEs, binding conditions that must read Yes before go-live, and each belongs to one of the four signers: compliance or legal, the clinical lead, the operations lead, and

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

the executive sponsor. Patients can decline the tool with their care unchanged: Yes. Patients are asked before any recording: Yes. A named change agent has the authority to slow or halt the go-live: Yes. But the fourth, a real person reviewing everything the tool creates before it enters a patient's record or is used for billing, reads Not yet. It is the same gap the approval carried forward, still open.

The regular items clear: staff have practiced in a safe test setting, and a named support resource is ready. One regular item reads Partially: the plain-language data-use explanation, waiting on the Portuguese materials.

The outcome. A STOPGATE at Not yet means the launch waits. The AI Governance Hub records the gap with an owner (the clinical lead) and a date, sets up the note-review step, and confirms that STOPGATE to Yes. It holds the data-use item until the Portuguese disclosure materials arrive and check out.

Then the AI Governance Hub selects “Conditional go-live (named gaps and dates)” and all four signers confirm the readiness area each one owns.

The 90-day pilot runs on schedule, and results are recorded in the plan's Living Record and compared against the Evidence Plan measures. Patient feedback during the pilot includes something staff had hoped for: patients notice their providers looking at them instead of the screen. When the tool moves from rollout to ongoing operations, the Stewardship Review Checklist (Section 6) takes over as the regular review instrument.

What this teaches. The two tools work as a chain. The plan carried every condition and gap forward, and the readiness checklist caught the open gap at the final step, before patients were involved. A STOPGATE that reads Not yet is the checklist doing its job. It turns an open gap into an owner, a date, and a launch the health center can stand behind. Bayside's story concludes at the end of Section 6.

REFERENCE APPENDICES

- Appendix A: SEEDS of Innovation™ — Full framework and maturity table for the five pillars used to measure the plan and confirm readiness in this section.
- Appendix B: Implementation Approach Reference — Framework descriptions and sourcing notes for the AI Implementation Readiness Checklist and the Playbook's other readiness tools.
- Appendix C: Glossary — Plain-language definitions for terms used in this section, including signal, go-live, STOPGATE, and Business Associate Agreement (BAA).
- Appendix D: Tools, Templates and Resources — The Design Plan Template and the AI Implementation Readiness Checklist, indexed with the Playbook's other tools.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Design Plan Template

This tool answers: How do we design a plan for implementation that works for our patients, our staff, and our health center?

Phase 3 — Plan and Implement (Section 5)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A template that turns an approved AI decision into an implementation plan, and a living record of how it performs.	Anchor the approved decision, build the plan (Part 2), run a pilot if needed, then review it (Part 3).	The AI Governance Hub files each plan with its Assessment and reviews the living record at every interval.

Before You Begin

This template picks up after a decision is made. It carries the approved decision from the AI Signal Decision Assessment (Section 4) into a concrete plan, then keeps a living record of how that plan performs.

Who fills this out: The AI Governance Hub members responsible for carrying out the decision, drawn from your Hub charter (Section 2). This is typically the Operations leadership, IT or Health Informatics, and Compliance or Legal perspectives, with Human Resources, Workforce, and Training and Communications as the plan reaches staff and patients. The executive sponsor (the leader who approves the plan) reviews the completed plan before it moves forward.

What you need first: A completed AI Signal Decision Assessment (Section 4) with an approved decision. If the decision has not been made yet, complete the AI Signal Decision Assessment first.

Part 1: Decision Reference

Part 1 anchors the plan to the decision that approved it. The full decision record lives in the AI Signal Decision Assessment (Section 4); transfer the key facts here so the plan and the decision stay connected.

IMPORTANT

This template is for decisions the AI Governance Hub approved or approved with conditions. If the recommendation was Hold or Do not proceed, that decision stays in the AI Signal Decision Assessment record. Return here only if the signal reopens and a new decision is approved.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI signal name	Name of the AI signal this plan addresses
Decision date	Date the decision was approved
Decision type	☐ New AI tool ☐ Change to existing tool ☐ Tool already in use without prior review
AI Signal Decision Assessment reference	Where to find the completed assessment (file location, system, or link)
Conditions or limits on the approval	Any conditions the AI Governance Hub placed on the decision. If none, write "None."
Unresolved gaps carried forward	Gaps the AI Governance Hub named but could not resolve during the assessment. These become items to address in this plan.

Part 2: The Plan

Part 2 is where the plan takes shape. Fill in what you know now, and return to update fields as decisions are made, timelines shift, or new information surfaces.

> **IMPORTANT**
> This plan is a living record. It is updated as your health center moves through implementation. Fill in what you know now and revisit each part as the work progresses.

Ownership

Assign each responsibility below to a named person on your AI Governance Hub, using the perspectives from your Hub charter (Section 2). These are the same roles the Hub already names; this plan does not create new ones. In a smaller health center, one person may hold several perspectives.

PERSPECTIVE (FROM YOUR HUB CHARTER)	NAME	RESPONSIBILITY IN THIS PLAN
Executive sponsor		Approves this plan, secures resources, and is accountable to leadership for whether the tool delivers value.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

PERSPECTIVE (FROM YOUR HUB CHARTER)	NAME	RESPONSIBILITY IN THIS PLAN
Operations leadership		Manages the rollout and how the tool fits into daily work; the first point of contact when something is not working.
IT or Health Informatics		System setup, data connections, security configuration, and ongoing technical maintenance.
Compliance or Legal		Confirms regulatory requirements, including HIPAA (the Health Insurance Portability and Accountability Act) and state privacy rules. Reviews the Business Associate Agreement (BAA) if one is required.
Human Resources, Workforce, and Training		Owns the workforce and workflow transition and the training plan.
Communications		Owns staff and patient communication, including plain-language AI disclosure.

Timeline

Set a target date and a responsible person for each milestone.

MILESTONE	TARGET DATE	RESPONSIBLE
Plan approved		
Technical setup complete		
Staff training complete		
Pilot start		
Pilot midpoint review		
Pilot end and go/no-go decision		
Full rollout (if approved)		

Training

Plan the training staff need before this tool goes live. Where the tool comes from a vendor, connect this to the vendor-provided training confirmed in the Vendor Evaluation Checklist (Section 4); the vendor should supply role-specific training materials.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Who needs training before go-live?	List the roles or teams
What do they need to learn?	What staff need to know or be able to do
How will training be delivered?	Format: in-person, virtual, self-paced, shadowing, etc.
How will you confirm staff are ready, and who confirms it?	How you will know training worked (observation, test run, sign-off), and the named person who confirms readiness

Workforce and Workflow Change

A new AI tool changes how people work. Name those changes before go-live so the plan supports staff through the transition, protects human oversight, and surfaces hidden or added work early rather than after launch.

What changes for staff	Which roles, tasks, or steps in the workflow change when this tool goes live
Workload effect	Whether the tool adds or removes steps, and where new manual or coordination work (for example, double-checking or reconciliation) may appear
Human oversight	Which decisions a person must review or be able to override, and who holds that responsibility
Support through the change	How staff will be supported through the transition: coaching, feedback, and an adjustment period
Who owns the workforce change	The Human Resources, Workforce, and Training perspective from your Hub charter (Section 2)

Communication Plan

Plan how staff, patients, and the community are told about this tool: who needs to know, what they need to understand, when, and how. Patients are told, in plain language, when AI is used in their care, and every audience is reached in the languages and formats they can understand.

Audience	Who needs to know: staff roles, patients, community partners
Message	What each audience needs to understand about this tool and how it affects them

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Timing	When each audience is told: before go-live, at launch, and on an ongoing basis
Channel	How the message reaches them: staff meeting, patient notice, signage, patient portal
Language and access needs	The languages and formats needed so every audience, including people who face the greatest barriers to access, can understand
Feedback route	How people ask questions or raise concerns, and who receives them
Owner and evidence of completion	The Communications perspective from your Hub charter (Section 2) owns each communication; note how you will confirm it happened

Evidence Plan

Track this tool against the same five SEEDS pillars the AI Signal Decision Assessment (Section 4) used to reach the decision. Define at least one measure for each pillar that applies; the Stewardship Review Checklist (Section 6) tracks these measures over time once the tool is live.

> **IN PLAIN TERMS**
> The five SEEDS pillars carry through every tool: Systemic Signal Assessment, Everyone-Centered Design, Ecosystem Alignment, Data Sovereignty, and Stewardship. Where a result could differ across the people you serve, set the measure by language or population group served, and name an example target so the goal is concrete.

SEEDS PILLAR	WHAT TO MEASURE	YOUR MEASURE(S) AND EXAMPLE TARGET
Systemic Signal Assessment	Whether the tool is addressing the real problem it was approved to solve.	Example: the workflow gap that prompted this tool is measurably smaller.
Everyone-Centered Design	Whether the patients and staff affected can access and safely benefit, including people who face the greatest barriers.	Example: no drop in visit access, and no rise in errors, by language or population group served.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

SEEDS PILLAR	WHAT TO MEASURE	YOUR MEASURE(S) AND EXAMPLE TARGET
Ecosystem Alignment	Whether the tool fits real workflows, staffing, and finances without adding hidden burden.	Example: note-completion time down 15%, with no rise in after-hours work.
Data Sovereignty	Whether data stays protected and is used only as approved.	Example: all required BAAs executed; no out-of-scope data use.
Stewardship	Whether the tool stays safe and accountable over time: reviews happen on schedule and issues are acted on.	Example: 100% of scheduled reviews completed on time; zero unresolved safety findings.

Guardrails

State plainly what this tool is and is not approved to do, and who must know those boundaries before rollout.

What this tool is approved to do	The specific scope of use. Be concrete: which patients, which workflows, which locations. Example: draft visit summaries for adult primary-care visits at the main site only.
What falls outside the approved scope	Boundaries. What this tool is not approved to do. Example: no use for pediatric visits or behavioral-health notes.
Who needs to know these boundaries	The roles or teams that must understand these guardrails before rollout begins
BAA (Business Associate Agreement) required?	☐ Yes ☐ No ☐ Under review
If yes: BAA status	☐ In place ☐ In negotiation ☐ Started
BAA includes AI-specific clauses?	☐ Yes ☐ No ☐ N/A. Look for clauses on data use for model training, data retention and deletion, secondary use, and liability if the tool causes harm.

Exit Criteria and Launch Guardrails

These are the conditions that must be true before this tool moves from pilot to full rollout. Complete this section during planning, then review it at the pilot midpoint and before the go/no-go decision. Every

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

item must be “Yes” (or “N/A”) before full rollout; if any item is “Not yet,” document what remains and set a date to revisit.

CRITERION	MET?
All success measures from the Evidence Plan have baseline data	☐ Yes ☐ Not yet
Staff training confirmed for all affected roles	☐ Yes ☐ Not yet
Technical validation complete (system tested in a production-like, or staging, environment)	☐ Yes ☐ Not yet
Monitoring and reporting systems active and producing data	☐ Yes ☐ Not yet
Shutdown protocol documented and tested	☐ Yes ☐ Not yet
No unresolved safety findings from the pilot	☐ Yes ☐ Not yet
Workforce and communication plans complete for all affected staff and patients	☐ Yes ☐ Not yet
BAA executed (if required)	☐ Yes ☐ Not yet ☐ N/A

IMPORTANT

When Part 2 is complete, review it with the executive sponsor for approval. If a pilot is planned, complete the Pilot Charter before starting. Then begin the Living Record (Part 3) to track how the plan performs.

Pilot Charter

Complete this section if your plan includes a pilot: a limited test before full rollout. The pilot charter defines what you are testing, how you will know whether it is working, and what would cause you to stop early.

Pilot Design

Define the shape of the pilot before it begins.

What are we testing?	The specific question this pilot answers
Where are we testing it?	Which site, department, or workflow

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Who is involved?	Which staff, and how many patients are affected
How long is the pilot?	Start date, end date, total duration
What do we predict will happen?	Your best expectation for results. This gives the midpoint review something to compare against.
Measures and targets for this pilot	From your Evidence Plan (the five SEEDS pillars), list the measures this pilot will test and the target each must reach to move to full rollout.

Midpoint Review

At the halfway point of the pilot, complete a Review Block (Part 3), comparing results to what you predicted and to your Evidence Plan targets. Decide whether to continue, adjust, or stop early. If you continue, move toward the pilot's go/no-go decision.

Early-Stop Triggers

These are conditions that would cause you to stop the pilot immediately, without waiting for the midpoint review. Define them before the pilot starts.

TRIGGER	THRESHOLD
Patient safety event	Any adverse event directly linked to the tool
Error rate	Set your threshold
Staff inability to use the tool	Set your threshold
System failure	Set your threshold
Add your own	

> **IMPORTANT**
> When an early-stop trigger fires: stop the pilot, document what happened, and bring the finding to the executive sponsor and the AI Governance Hub before deciding next steps.

Part 3: Living Record

The Living Record tracks how this plan performs over time, using the same Review Block at every review point.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

The Review Block

Use one Review Block at every review point: the pilot midpoint (if you ran a pilot) and each post-rollout review. The first post-rollout review is typically 30, 60, or 90 days after go-live, depending on the tool's risk level and your capacity. Copy the block for each round and number them (Round 2, Round 3, and so on); the Living Record stays active for the full life of this tool's implementation phase.

Review date and reviewer(s)	
Stage	☐ Pilot midpoint ☐ Post-rollout review — Round ____
What the evidence shows	The data, reports, and feedback reviewed for this review
How it compares to the plan's targets	Compare to your Evidence Plan targets — and, at the pilot midpoint, to what you predicted
What is working	
What needs adjustment	
Decision	☐ Continue as planned ☐ Continue with adjustments ☐ Pause for further review ☐ Stop or begin exit planning
If adjusting: what changes, and who is responsible	
Next review date	

Transition to Stewardship

When this tool moves from initial rollout to ongoing operations, the Stewardship Review Checklist (Section 6) takes over as the regular review instrument. Schedule the first stewardship review before closing the final round of this Living Record.

Date this tool moved to ongoing operations	
First Stewardship Review Checklist scheduled for	
Handed off by	Name of the person closing the Living Record

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Plan Quality Check

Before this plan moves forward, walk through each item below. Every item should be checked. If any item cannot be checked, return to the relevant section and complete it.

#	CHECK	✓
1	The Decision Reference (Part 1) connects this plan to an approved AI Signal Decision Assessment, including any conditions and unresolved gaps carried forward.	☐
2	Every ownership role in Part 2 has a named person.	☐
3	The timeline has target dates for each milestone.	☐
4	Training is planned: who needs training, what they need to learn, and who confirms they are ready.	☐
5	Workforce and workflow changes are named, including human-oversight responsibilities.	☐
6	A communication plan is in place for staff and patients, with language and access needs addressed.	☐
7	The Evidence Plan has at least one measure and target for each SEEDS pillar that applies.	☐
8	Guardrails state what this tool is and is not approved to do, and who must know them.	☐
9	BAA status is documented (if applicable).	☐
10	Exit criteria are defined for moving from pilot to full rollout.	☐
11	If a pilot is planned, the Pilot Charter is complete, including its measures, targets, and early-stop triggers.	☐
12	The people affected by this tool (patients, staff, or both) are identified and their needs are reflected in the plan.	☐
13	The first Living Record review date is set.	☐
14	The executive sponsor has reviewed and approved this plan.	☐

Approval note: the AI Governance Hub approved the decision (recorded in the AI Signal Decision Assessment). The executive sponsor approves this implementation plan. The AI Governance Hub then files the approved plan and tracks the Living Record.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Plan approved by	Executive sponsor name
Date approved	
Plan filed with the AI Governance Hub	☐ Yes

KEY TAKEAWAYS BY ROLE

Leadership

- This is the plan that determines whether an AI decision succeeds or stalls. Review it before approving, and hold the team accountable for the milestones, guardrails, and review schedule it sets.
- The Living Record and the Stewardship Review Checklist (Section 6) are how you will know whether this tool is delivering what it was approved to do.

AI Governance Hub members

- File each completed Design Plan Template alongside the AI Signal Decision Assessment for that signal.
- Track the Living Record review dates and follow up when reviews are due.
- When the tool transitions to ongoing operations, schedule the first Stewardship Review Checklist.

IT & Security

- Own the technical fields: setup timeline, data connections, BAA status, AI-specific contract clauses, and the shutdown protocol in the exit criteria.
- Flag technical risks early in planning, before the pilot starts, rather than during it.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**
This template comes after a decision. Start at Phase 1 (Section 3) to catch and assess the signal.

→ **Evaluating a specific AI tool?**
Complete the AI Signal Decision Assessment (Section 4) first. This template picks up after the decision is approved.

→ **Formalizing an existing AI tool?**
Complete the AI Signal Decision Assessment to document the decision, then use this template to build the plan for keeping it accountable.

→ **Reassessing a changing AI tool?**
If the change reopens the decision, update the AI Signal Decision Assessment first, then revise this plan to match.

→ **Replacing or exiting an AI tool?**

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Begin exit planning in the Living Record decision, and hand off to the Stewardship Review Checklist (Section 6).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI Implementation Readiness Checklist

This tool answers: Is our organization ready to launch this?

Phase 3 — Plan and Implement (Section 5)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
A go-live checklist that confirms people, patients, and safeguards are ready before an approved AI tool launches.	Work each pillar with your team before launch, marking Yes, Partially, or Not yet. Every STOPGATE must read Yes.	The AI Governance Hub confirms readiness and signs the go-live decision, keeping the completed checklist on file.

How to use this checklist

This checklist confirms the last mile before an approved AI tool launches: that people and patients are ready, and that the launch-day safeguards are live.

The Design Plan (Section 5) builds and gates the plan. This checklist verifies what that plan set in motion, so each question confirms readiness rather than re-collecting what the Design Plan already records.

The checklist opens at Everyone-Centered Design. The systemic problem and scope were settled upstream in the AI Signal Decision Assessment (Section 4) and the Design Plan, so this tool begins with the people and patients the launch touches.

> **IN PLAIN TERMS**
> Go-live means the day the tool officially starts being used in real work, with real patients or staff.

Reading the Scope Tags

Each question ends with one or more tags.

Acquisition — which path the tool came from:

- **[Buy]** applies when the health center purchased the AI tool from a vendor.
- **[Build]** applies when the health center developed the tool internally.
- **[Buy or Build]** applies in every case.

Context — where the tool operates:

- **[Clinical]** applies when the AI tool supports direct patient care or clinical decisions.
- **[Admin]** applies when the AI tool supports administrative or operational functions.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- No context tag means the question applies in both settings.

Binding condition:

- **STOPGATE** halts the go-live process until it is resolved, because going forward would create harm or exposure that cannot be undone after the fact.

Bracketed source keys identify the framework each question draws on; full citations appear in Appendix B: Implementation Approach Reference.

Pillar 2: Everyone-Centered Design

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P2.Q1 Patients can decline this tool in their care, and their care stays exactly the same. **[Buy or Build]** **[Clinical]** **STOPGATE** *[WHO] [SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P2.Q2 The patient advisory board or committee has been briefed before go-live. **[Buy or Build]** *[HRSA 330] [Betsy Lehman]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 3: Ecosystem Alignment

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P3.Q1 Frontline staff know the tool is coming and understand what it does and how it helps their work. **[Buy or Build]** *[NCFH] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q2 Staff have had hands-on practice in a non-production setting before go-live. **[Buy or Build]** *[HSCC] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q3 Change agents (project managers, department heads, clinical leads, or AI champions) can explain the rollout to the teams they lead. **[Buy or Build]** *[NCFH]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q4 A named support resource is available to staff for questions once the tool is live. **[Buy or Build]**	☐ Yes ☐ Partially	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
[NCFH] [DiMe]	☐ Not yet	
P3.Q5 Staff know what to do if the tool is unavailable or produces an unexpected result. **[Buy or Build]** *[HSCC] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	
P3.Q6 A real person reviews anything the tool creates before it goes into a patient's record or is used for billing, and that review is operating before launch. **[Buy or Build]** **STOPGATE** *[HSCC]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 4: Data Sovereignty

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P4.Q1 For tools that listen or record during a visit, patients are asked beforehand whether they are comfortable with it; they can say no, and their care stays the same. **[Buy or Build]** **[Clinical]** **STOPGATE** *[HSCC] [SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q2 Patients receive a plain-language explanation of how their data is used by this tool, covering what is collected, who can access it, how long it is kept, and how it is protected, before go-live. **[Buy or Build]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q3 A named person or body holds the authority to refuse uses of patient data outside the approved scope, such as vendor model retraining or resale, and that boundary is set before go-live. **[Buy]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	
P4.Q4 Staff who handle the tool can explain, in plain terms, where patient data goes and who has access. **[Buy or Build]** *[SEEDS]*	☐ Yes ☐ Partially ☐ Not yet	

Pillar 5: Stewardship

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
P5.Q1 A named change agent has the authority, and a channel to the AI Governance Hub, to slow or halt the go-live if frontline use shows a problem. **[Buy or Build]** STOPGATE *[NIST AI RMF] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	
P5.Q2 Escalation paths are confirmed for problems that surface at or after launch. **[Buy or Build]** *[HSCC] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	

IMPORTANT

If your team cannot answer a question at all, especially in the staff-readiness items, that is the clearest sign that more preparation is needed before the tool starts. A Not yet you can see is one you can act on.

Binding conditions (STOPGATEs)

Each binding condition must read Yes before go-live, and each is owned by one signer below: (P4.Q1, Compliance or Legal); (P3.Q6, Clinical lead); (P5.Q1, Operations or Implementation lead); (P2.Q1, Executive sponsor).

Go-Live Decision and Sign-Off

Go-live proceeds when every binding condition reads Yes and each signer confirms the readiness area they own. Where a gap remains, record it with an owner and a date, and choose a conditional path.

Decision: [ ] Go live [ ] Conditional go-live (named gaps and dates) [ ] Hold

ROLE	PRINT NAME	SIGNATURE	DATE
Compliance or Legal			
Clinical lead			
Operations or Implementation lead			
Executive sponsor			

KEY TAKEAWAYS BY ROLE

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Leadership	AI Governance Hub members	IT & Security
• Confirm every binding condition reads Yes before you authorize go-live. • You own the patient opt-out condition: patients can decline the tool with their care unchanged.	• Run this checklist with the implementation team before launch, and keep it on file with the Design Plan. • Route any Not yet to an owner with a date before the go-live decision.	• Confirm the live human-review step for any output entering the record or billing. • Confirm audit and fallback paths so staff know what to do if the tool is unavailable.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**
Start at Phase 1 — Listen and Learn (Section 3). This checklist runs much later, just before go-live.

→ **Evaluating a specific AI tool?**
Assess it first with the AI Signal Decision Assessment (Section 4), then plan it with the Design Plan (Section 5). This checklist confirms readiness once the plan is built.

→ **Formalizing an existing AI tool?**
Document the decision in the AI Signal Decision Assessment (Section 4) and build the Design Plan (Section 5); return here before it goes live.

→ **Reassessing a changing AI tool?**
If the change reopens the decision, update the AI Signal Decision Assessment (Section 4) and the Design Plan (Section 5), then re-run this checklist before relaunch.

→ **Checking an approved AI tool?**
Use the Stewardship Review Checklist (Phase 4 — Steward, Section 6) for tools already live. This checklist applies at the moment of launch.

→ **Replacing or exiting an AI tool?**
Plan the change through the AI Signal Decision Assessment (Section 4) and the Design Plan (Section 5); return here before the replacement goes live.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 6: Phase 4 — Steward

GOAL OF THIS SECTION

Keep every live AI tool accountable: measure its performance each cycle, decide whether to continue, adjust, pause, or exit, and share what you learn.

WHAT'S COVERED IN THIS SECTION

6.1 Measuring Performance and Impact
Score five SEEDS pillars against Green, Amber, Red thresholds.

6.2 Adjust, Scale, or Stop
Turn findings into a continue, adjust, pause, or exit decision.

6.3 Feeding Stewardship Back to the AI Governance Hub
Route findings back into governance and network learning.

TOOLS YOU'LL USE

Stewardship Review Checklist
Is this tool still delivering what we said it would — what do the measures show, and what do we do about it?

Case Example Sharing Template
How do we share what we learned?

6.1 Measuring Performance and Impact

When you review a live AI tool, you look at the same five SEEDS pillars that guided the decision to adopt it. Each pillar asks one plain question: is this tool still doing what we said it would?

The **Stewardship Review Checklist** organizes every review around these five pillars. Part A (Performance Measures) sets a Green, Amber, and Red threshold for each measure, so your AI Governance Hub can see at a glance where a tool stands and where it needs attention. Part B (Review and Decision) marks each pillar item Yes, Partially, or Not yet; items marked Partially or Not yet become the action list for the next cycle.

Before each review cycle, confirm the right evidence is in hand for each pillar. The table below maps each pillar to the accountability question it answers and the evidence your AI Governance Hub gathers. The AI Signal Decision Assessment (the full review your health center runs before approving a tool; Section 4: Phase 2 — Assess and Decide) set the original expectations; the stewardship review checks whether they held.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

SEEDS PILLAR	ACCOUNTABILITY QUESTION	EVIDENCE THE HUB GATHERS
1. Systemic Signal Assessment	Is the tool still meeting the need it was adopted to solve?	Intended-use adherence against the Approved-Use Inventory Checklist; whether the tool still supports current organizational priorities
2. Everyone-Centered Design	Is the tool serving every patient well?	Patient-reported experience (existing surveys or the Patient AI Feedback Form Template); outcome consistency across the populations served; access impact
3. Ecosystem Alignment	Is the tool working reliably at a sustainable cost?	Accuracy against baseline; reliability and integration stability; workflow fit and staff capacity; total cost of ownership against budget; UDS clinical quality impact
4. Data Sovereignty	Is patient data protected and the model still valid?	Security and access events; vendor Business Associate Agreement status; data or model drift
5. Stewardship	Are we watching for harm and able to act?	Adverse-event reporting; a tested shutdown and rollback path; whether the review cadence held

IN PLAIN TERMS

The Uniform Data System (UDS) is the federal reporting system every Federally Qualified Health Center uses to show how it serves its community. If you already report UDS clinical quality measures, you have a foundation for tracking a tool's impact on the patients you serve.

IMPORTANT

If your health center did not capture baseline data before the tool went live, start collecting it now. A review without a baseline is still worth doing: use the first review to set the baseline for every review that follows.

6.2 Adjust, Scale, or Stop

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Part A gives every measure a threshold: Green (performing as expected), Amber (needs attention within the next cycle), and Red (requires a formal Stewardship Review before continued use or expansion). Before any decision, the AI Governance Hub reads all the evidence against these thresholds. A Red on patient safety or data protection pauses the tool even when every other measure is strong; a Red overrides the overall picture.

The response is graduated by threshold:

- **Green — Continue.** Record the status in the Review Cycle Log and confirm the next review date.
- **Amber — Focus.** The pillar's subject-matter lead investigates and reports back at the next cycle. Expansion resumes when Amber returns to Green.
- **Red — Stop and review.** Pause expansion. The AI Governance Hub Lead convenes a formal Stewardship Review (Part B) within 30 days.

Part B, section B.4, records the Hub's decision. Four options are available:

- **Continue as is** — the evidence supports keeping the tool in place. The Hub schedules the next review and logs the decision.
- **Continue with adjustments** — one or more pillars show a gap. The Hub names the issue, assigns an owner, and sets a timeline; the gap is resolved before the next review.
- **Pause for further review** — the evidence needs closer examination before the tool continues or expands. The Hub sets a timeline and names the additional data needed.
- **Begin exit planning** — a threshold has been crossed, or the evidence no longer supports continuing. Part B's decommissioning checklist (B.5) walks the Hub through data handling, vendor notification, fallback activation, and staff communication.

> **IMPORTANT**
> Part B includes a policy and boundary check. A “Yes” to any item — the tool used outside its approved scope, an unresolved patient safety concern, a vendor change to terms or data handling, or a regulatory change affecting use — escalates to the AI Governance Hub before the tool continues, whatever the pillar scores show.

6.3 Feeding Stewardship Back to the AI Governance Hub

Stewardship findings move back into governance through three pathways.

Standing review schedule. Stewardship Review Checklist findings from each active tool feed the next AI Governance Hub meeting agenda. The Hub reviews what changed, confirms what decisions are needed, and records what to update: the Approved-Use Inventory, the monitoring frequency, or the AI policy if practice has drifted from what the policy describes. The Review Cycle Log in Part A is the audit trail — each entry records when the review happened, what the Hub found, and what it decided.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Re-entry into the governance cycle. Some findings restart the governance cycle for a tool. A finding that crosses a Red threshold opens the formal Stewardship Review (Part B); a reopened decision goes back through Phase 2 — Assess and Decide (Section 4) to revisit the decision and the Pre-Deployment Impact Checklist assumptions, then through Phase 3 — Plan and Implement (Section 5) to update the plan and close the record. A pattern across several tools — the same vendor issue in two places, the same workflow burden in different departments — may open a broader review through the AI Signal Decision Assessment (Section 4: Phase 2 — Assess and Decide).

Network learning. Findings with value beyond your health center move to the Massachusetts League of Community Health Centers through the **Case Example Sharing Template**. A signal (any indication from a person, system, vendor, regulator, or peer organization that something AI-related needs attention) that a review surfaces routes through the AI Signal Intake Form Template so the Hub can track it. The sharing template welcomes four types of case example: a tool that worked; a tool that did not work and what the Hub learned; a decision to wait and the reasoning behind it; and AI already in use and what the Hub did to bring it into governance. Case examples are anonymized by default, and attribution is opt-in. A health center that chooses to share emails the completed template to the Massachusetts League at MA.HCCN@massleague.org. The Massachusetts League uses what health centers share to surface patterns, inform Playbook updates, and issue peer guidance; your Health Center Controlled Network (HCCN) is another channel for sharing findings with peers.

CASE EXAMPLE: THE ZOMBIE PILOT PATTERN AT BAYSIDE

This example is hypothetical and depicts a composite community health center. It concludes the Bayside story.

The story so far. Bayside Community Health Center built its **AI Governance Hub** (Section 2), listened to staff and patients (Section 3), approved an ambient documentation tool with conditions (Section 4), and launched a 90-day pilot that met its goals (Section 5).

The situation. Bayside did everything right through launch. Then time passed. The pilot succeeded, the tool expanded to more providers, and the quarterly review schedule quietly slipped. A year after go-live, two providers at a second site started using the tool without AI Governance Hub approval, because it was already working well at the first site.

Nobody decided to expand. It just happened. This is the zombie pilot pattern: a tool that keeps spreading because no one scheduled the moment to judge it.

The process. The AI Governance Hub Coordinator notices the missed reviews and puts the tool on the next agenda. The AI Governance Hub runs the overdue review using the Stewardship Review Checklist. Part A scores each pillar's measures against set Green, Amber, and Red thresholds.

Systemic Signal Assessment comes back Red: the tool is being used outside its approved scope. Stewardship comes back Red: two review cycles were missed. Everyone-Centered Design comes back Amber: patient disclosure at the second site never happened, because that site was never formally added.

A Red score does two things by itself: expansion pauses, and a formal Stewardship Review (Part B) opens within 30 days. In Part B's policy and boundary check, the question “Has the tool been used

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

outside its original approved scope?" gets a Yes, which escalates to the full AI Governance Hub before the tool can continue.

Then the AI Governance Hub selects from Part B's four decision options: "Continue with adjustments." The expansion stays paused. The disclosure process reaches every patient at both sites. A quarterly review schedule is set again, this time with calendar invitations and a named owner. Only then does continued use get approved.

The review's learning-capture question asks whether anything is worth sharing with other health centers. The answer is yes, so the AI Governance Hub completes a Case Example Sharing Template, marks the type "AI already in use, and what the AI Governance Hub did to bring it into governance," and submits it anonymously to the Massachusetts League at MA.HCCN@massleague.org.

The outcome. The tool stays in use, now on a real review schedule. Disclosure reaches every patient. Expansion resumes only after the next review returns the measures to Green. And the whole network gains a documented example of how a well-run launch can still drift, and how a health center brings it back.

What this teaches. Good governance at launch is never the finish line. A tool without a scheduled review date spreads by habit instead of by decision, even at a health center that did everything else right. The Stewardship Review Checklist catches the pattern because it scores governance alongside performance. The Case Example Sharing Template turns one health center's course correction into learning every health center can use.

The tools in this section are indexed in Appendix D: Tools, Templates and Resources; the frameworks behind the Part A measures are cited in Appendix B: Implementation Approach Reference; the SEEDS pillar definitions are in Appendix A: SEEDS of Innovation™.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Stewardship Review Checklist

This tool answers: Is this tool still delivering what we said it would — what do the measures show, and what do we do about it?

Phase 4 — Steward (Section 6)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
Part A measures each active AI tool against Green/Amber/Red thresholds; Part B reviews, decides, and records the outcome.	Score Part A measures each cycle; any Red opens Part B: the five-pillar review, policy check, and decision.	The Hub runs each review, owns the decision, and keeps the completed checklist in the tool's governance record.

Reading the Scope Tags

Each question ends with one or more tags.

Acquisition — which path the tool came from:

- **[Buy]** applies when the health center purchased the AI tool from a vendor.
- **[Build]** applies when the health center developed the tool internally.
- **[Buy or Build]** applies in every case.

Context — where the tool operates:

- **[Clinical]** applies when the AI tool supports direct patient care or clinical decisions.
- **[Admin]** applies when the AI tool supports administrative or operational functions.
- No context tag means the question applies in both settings.

Binding condition:

- **In this checklist, a STOPGATE is a policy and boundary item in Part B.3; a Yes answer escalates to the AI Governance Hub before the tool continues in use.**

Bracketed source keys identify the framework each question draws on; full citations appear in Appendix B: Implementation Approach Reference.

IN PLAIN TERMS

The Uniform Data System (UDS) is the federal reporting system every Federally Qualified Health Center uses to show how it serves its community. If you already report UDS clinical quality measures, you already have a foundation for tracking a tool's impact on the patients you serve.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Part A — Performance Measures

Review Part A each cycle (quarterly is a good starting point). The AI Governance Hub Lead assigns each pillar's accountability question to a subject-matter lead, who gathers the evidence before the Hub meets. For each measure, match your data to the Green, Amber, or Red threshold, and skip any measure whose scope tag does not match your tool. Record every finding in the Review Cycle Log. Amber opens a focused review by the pillar lead; Red opens a formal Stewardship Review (Part B) within 30 days, and expansion pauses until it resolves.

P1 · Systemic Signal Assessment

Accountability question: Is the tool still meeting the need it was adopted to solve?

MEASURE	GREEN	AMBER	RED
P1.Q1 Intended-use adherence: share of tool usage within the approved scope on the Approved-Use Inventory Checklist. **[Buy or Build]** *[HSCC] [SEEDS]*	95%+ within approved scope	85–94% within scope	Below 85% within scope
P1.Q2 Strategic alignment: the tool still supports current organizational priorities and the community need it was adopted to address. **[Buy or Build]** *[Balanced Scorecard] [SEEDS]*	Aligned with current priorities	Partial; priorities shifting	No longer supports current direction

P2 · Everyone-Centered Design

Accountability question: Is the tool serving every patient well?

MEASURE	GREEN	AMBER	RED
P2.Q1 Patient-reported experience: feedback about interactions involving the tool (existing surveys, comment cards, or the Patient AI Feedback Form Template). **[Buy or Build]** *[Betsy Lehman]*	No AI-related concerns	Isolated concerns; no pattern	Pattern of concerns or any report of harm
P2.Q2 Outcome consistency across patient groups: outcomes compared across the populations served (age, language, insurance, geography). **[Buy or Build]** *[DiMe] [WHO]*	No meaningful differences	Differences under investigation	Confirmed outcome gap linked to the tool

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

MEASURE	GREEN	AMBER	RED
P2.Q3 Access impact: whether the tool changed how patients reach care (availability, wait times, communication). **[Buy or Build]** *[Donabedian] [NACHC VTF]*	Access stable or improved	Access declined for a group	Measurable access barrier created

P3 · Ecosystem Alignment

Accountability question: Is the tool working reliably at a sustainable cost?

MEASURE	GREEN	AMBER	RED
P3.Q1 Accuracy versus baseline: share of outputs matching expected results against the baseline set at deployment. **[Buy or Build]** *[DiMe] [Donabedian]*	Meets or exceeds baseline	5–10% decline	More than 10% decline
P3.Q2 Reliability: system uptime against the SLA or internal target, and integration stability at the EHR connection. **[Buy or Build]** *[HSCC] [ONC SAFER]*	Meets SLA; rare integration errors	One SLA miss or intermittent errors	Two+ misses or frequent/data-sync failures
P3.Q3 Workflow fit and staff capacity: override rate and staff-reported ease of use (brief pulse survey). **[Buy or Build]** *[ONC SAFER] [NCFH]*	Override stable; staff report good fit	Rising override or mixed feedback	Override >25% or majority report disruption
P3.Q4 Total cost of ownership: all tool costs (licensing, maintenance, staff time, infrastructure) against the budget approved at deployment. **[Buy or Build]** *[DiMe] [Balanced Scorecard]*	Within approved budget	5–15% over budget	More than 15% over budget
P3.Q5 UDS clinical quality impact: change in the Uniform Data System clinical quality measures the tool is expected to influence. **[Buy or Build]** **[Clinical]** *[NACHC VTF] [Donabedian]*	Quality measures stable or improving	Any measure declined since deployment	Measure below prior-year benchmark

P4 · Data Sovereignty

Accountability question: Is patient data protected and the model still valid?

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

MEASURE	GREEN	AMBER	RED
P4.Q1 Security and access control: security events involving the tool, and the share of users with appropriate, current access. **[Buy or Build]** *[HSCC] [NIST AI RMF]*	Zero incidents; access current	One contained event, no PHI; 1–3 access gaps	Any PHI exposure or 4+ access gaps
P4.Q2 Vendor security and BAA: vendor maintains a current Business Associate Agreement and meets contractual security terms. **[Buy]** *[HSCC]*	BAA current; terms met	One term under discussion	BAA lapsed or a security term failed
P4.Q3 Data or model drift: evidence that outputs have shifted from validated performance, within accepted tolerance (machine-learning tools). **[Buy or Build]** *[DiMe] [ECRI]*	No drift or within tolerance	Drift detected; investigating	Drift confirmed with impact

P5 · Stewardship

Accountability question: Are we watching for harm and able to act?

MEASURE	GREEN	AMBER	RED
P5.Q1 Adverse-event reporting active: adverse events, near-misses, and unexpected results are captured and reviewed through existing reporting systems. **[Buy or Build]** *[ECRI] [CHAI]*	Reporting active; none unresolved	An event under review	Reporting inactive or an unresolved harmful event
P5.Q2 Shutdown and rollback tested: a manual or automated way to disengage the tool is tested and functional. **[Buy or Build]** *[NIST AI RMF] [HSCC]*	Tested and functional	Exists; not recently tested	Absent or fails testing
P5.Q3 Review cadence held: the tool has been reviewed on its set schedule, with findings recorded. **[Buy or Build]** *[NIST AI RMF] [IHI PDSA]*	Reviewed on schedule	One cycle missed	Two+ cycles missed

Review Cycle Log

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Record each review cycle. The log is the audit trail of when the tool was reviewed, what was found, and what was decided. The AI Governance Hub Lead maintains it.

REVIEW DATE	P1	P2	P3	P4	P5	ACTION TAKEN	NEXT REVIEW

> **IMPORTANT**
> Beyond the thresholds, open a formal Stewardship Review right away for any of these: a patient safety event, a regulatory change affecting use, a vendor breach or BAA lapse, a whistleblower concern submitted through the AI Signal Intake Form Template, or a payer notification or audit inquiry.

Part B — Review and Decision

Complete Part B when any Part A measure reads Red, or on the tool's scheduled full review. Mark each five-pillar item Yes, Partially, or Not yet based on the evidence; complete the policy and boundary check; then record the outcome, decision, and learning together.

B.1 Cover sheet

Tool name	
Review date	
Reviewer(s)	
Review type	☐ Scheduled ☐ Triggered by incident ☐ Triggered by change
Last review date	
Review cadence	☐ Quarterly ☐ Semi-annual ☐ Annual ☐ Other: ____

B.2 Five-pillar review

Mark each item Yes, Partially, or Not yet. Use the third column to note the evidence reviewed and any follow-up for a Partially or Not yet.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

QUESTION	YES / PARTIALLY / NOT YET	EVIDENCE / SOURCE
B2.Q1 Does the tool still align with the health center's mission, strategic plan, and current priorities? **[Buy or Build]** *[Balanced Scorecard] [Kotter]*	☐ Yes ☐ Partially ☐ Not yet	
B2.Q2 Has patient or community feedback been collected since the last review and reviewed across the populations served? **[Buy or Build]** *[Betsy Lehman] [NACHC VTF]*	☐ Yes ☐ Partially ☐ Not yet	
B2.Q3 Are the right people still in the roles responsible for this tool, and is staff capacity holding? **[Buy or Build]** *[NIST AI RMF] [ONC SAFER]*	☐ Yes ☐ Partially ☐ Not yet	
B2.Q4 Are safety and security monitoring systems active and producing regular reports? **[Buy or Build]** *[ONC SAFER] [DiMe]*	☐ Yes ☐ Partially ☐ Not yet	
B2.Q5 Is the AI Governance Hub reviewing this tool on its set cadence, and is the shutdown protocol functional? **[Buy or Build]** *[NIST AI RMF] [ECRI]*	☐ Yes ☐ Partially ☐ Not yet	

B.3 Policy and boundary check

These are binding conditions. A "Yes" to any item triggers immediate escalation to the AI Governance Hub before the tool continues in use, regardless of the five-pillar results.

POLICY CHECK (BINDING CONDITION)	YES	NO	NOTES
B3.Q1 Has the tool been used outside its original approved scope? **[Buy or Build]** **STOPGATE** *[HSCC] [SEEDS]*	☐ Yes	☐ No	
B3.Q2 Has a patient safety concern been raised that remains unresolved? **[Buy or Build]** **STOPGATE** *[ECRI] [CHAI]*	☐ Yes	☐ No	
B3.Q3 Has the vendor changed terms, pricing, or data handling since the last review? **[Buy]** **STOPGATE** *[HSCC]*	☐ Yes	☐ No	

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

POLICY CHECK (BINDING CONDITION)	YES	NO	NOTES
B3.Q4 Has a regulatory or compliance change affected how this tool can be used? **[Buy or Build]** STOPGATE *[HSCC] [NIST AI RMF]*	☐ Yes	☐ No	

IMPORTANT

If any policy item is "Yes," document the issue and escalate to the AI Governance Hub before completing the decision below.

B.4 Outcome, decision, and learning

Record the review result, the decision, and what the team learned in one place.

Review result P1 ☐ / P2 ☐ / P3 ☐ / P4 ☐ / P5 ☐ (Yes · Partially · Not yet per pillar) **Stopgates triggered: ☐ Yes ☐ No**

DECISION	SELECT	FOLLOW-UP ACTIONS	TIMELINE
Continue as is	☐		
Continue with adjustments	☐		
Pause for further review	☐		
Begin exit planning	☐		

A triggered stopgate in B.3 moves the tool to the AI Governance Hub for escalation review regardless of the pillar results.

Each binding condition is owned by one signer below: (B3.Q1, Executive sponsor); (B3.Q2, Clinical lead); (B3.Q3, IT or Security lead); (B3.Q4, Compliance or Legal). The decision proceeds when each signer confirms the conditions they own.

ROLE	PRINT NAME	SIGNATURE	DATE
Executive sponsor			
Clinical lead			
IT or Security lead			
Compliance or Legal			

Learning capture

Capture the lessons from this review. If the last answer is yes, complete a Case Example Sharing Template (Section 6) so other health centers can learn from it.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

What worked well since the last review?	
What surprised us?	
What would we do differently?	
Anything worth sharing with other health centers?	

B.5 Decommissioning checklist

Complete this section only if the decision is "Begin exit planning."

STEP	DONE	NOTES
Notify all staff currently using the tool	☐	
Identify and secure any patient data stored in the tool	☐	
Confirm the data export or migration plan with the vendor	☐	
Review contract exit terms (notice period, data return, fees)	☐	
Plan workflow changes for the transition period	☐	
Communicate the change to patients if the tool was patient-facing	☐	
Document lessons learned in the learning capture above	☐	
Update the AI Governance Hub tool inventory	☐	
Set a follow-up date to confirm decommissioning is complete	☐	

KEY TAKEAWAYS BY ROLE

Leadership

- Own the decision to act on Red findings; a Red pauses expansion until the Stewardship Review completes.
- Set the review cadence and confirm it at least once a year; quarterly is a good starting point.

AI Governance Hub members

- Assign each pillar to a subject-matter lead before each cycle; that person gathers the evidence.
- Track Amber to resolution; a persistent Amber escalates to Red.

IT & Security

- Provide the Data Sovereignty evidence: security events, access audits, and drift monitoring.
- Confirm the shutdown and rollback path is tested and functional before the review.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

• Use the decision and sign-off as the board-reportable record of accountability.	• A triggered stopgate moves the tool to escalation regardless of the pillar results.	• Flag logging or monitoring gaps before the review, not during it.

WHERE THIS FITS

→ **Exploring whether an AI tool is the right fit?**

Start at Phase 1 — Listen and Learn (Section 3). This checklist applies after a tool is live and ready for its first review.

→ **Evaluating a specific AI tool?**

Assess it with the AI Signal Decision Assessment (Section 4), then plan it with the Design Plan (Section 5). Return here once the tool is live.

→ **Formalizing an existing AI tool?**

Bring it into governance through Phase 2 — Assess and Decide (Section 4); once it is live, use this checklist as its first baseline review.

→ **Reassessing a changing AI tool?**

Run this review as an unscheduled check. The policy and boundary check (Part B) is built for a vendor change, scope expansion, or regulatory shift.

→ **Checking an approved AI tool?**

You are here. Score Part A each cycle and record findings in the Review Cycle Log; any Red opens Part B.

→ **Replacing or exiting an AI tool?**

Complete this review to document the current state, then the decommissioning checklist (Part B). For the replacement, start at Phase 1 — Listen and Learn (Section 3).

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Case Example Sharing Template

This tool answers: How do we share what we learned?

Phase 4 — Steward (Section 6)

WHAT THIS TOOL IS	HOW THIS TOOL WORKS	AI GOVERNANCE HUB APPLICATION
Documents one AI governance case example: what the AI Governance Hub decided, what happened, and what the network can learn.	Complete it after a governance cycle reaches a clear outcome. Answer each section in plain language, in your own words.	The AI Governance Hub reviews each completed template and decides which case examples to share with the Massachusetts League.

Before You Begin

The AI Governance Hub Lead or a designated AI Governance Hub member completes this template. The AI Governance Hub reviews every submission before it is shared.

Every case example strengthens the network. This template welcomes four types of case examples:

- A tool that worked, and what made it work
- A tool that did not work, and what the AI Governance Hub learned
- A decision to wait, and the reasoning behind it
- AI already in use, and what the AI Governance Hub did to bring it into governance

A case example can center on what the AI Governance Hub learned about implementation and change management, such as how staff adopted the tool, what the rollout surfaced, and what the Hub adjusted, not only whether the tool itself worked.

This template captures what your health center learned from governing, implementing, and managing change around AI, so the network can learn from it.

Sharing and Privacy

Case examples are anonymized by default. Each submitting health center chooses whether to attach its name, and attribution is opt-in.

☐ Anonymous. Share this case example without identifying our health center.
☐ Attributed. We give the Massachusetts League permission to identify our health center by name.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Patient names and identifying details are always excluded. Aggregate, de-identified data, such as accuracy rates by language group over a six-month period, is appropriate and often the most useful evidence for network learning.

Staff and clinician names are replaced with roles, such as the AI Governance Hub Lead or the AI Tool Owner. The Massachusetts League reviews submissions for network distribution and does not edit substantive content. The submitting health center retains authorship.

The Template Starts Here

Answer each question below in your own words. The italic examples come from one hypothetical case example, for reference only. Where a question does not apply, write "N/A" and move on.

What type of case example is this? (Check one.)

☐ A tool that worked, and what made it work
☐ A tool that did not work, and what the AI Governance Hub learned
☐ A decision to wait, and the reasoning behind it
☐ AI already in use, and what the AI Governance Hub did to bring it into governance

One-sentence summary:

Example: Our AI Governance Hub approved ambient documentation for primary care visits, and after six months the tool met its goals for provider time savings and note quality, with adjustments needed for patient disclosure and multilingual accuracy.

Three-sentence overview:

Example: A mid-size community health center adopted an ambient documentation tool to reduce after-hours charting for primary care providers. The AI Governance Hub used the AI Signal Decision Assessment and Pre-Deployment Impact Checklist to evaluate it, approved a rollout with three providers, and expanded to twelve after the first stewardship review. The network can learn from how the Hub structured its disclosure process and caught a documentation gap affecting patients whose primary language was Portuguese.

What type of AI tool is involved? (Check all that apply.)

☐ Ambient documentation (automatic note generation during visits)
☐ Clinical decision support
☐ Revenue cycle or billing automation
☐ Patient communication (chatbot, portal messaging)
☐ Scheduling or call center automation
☐ Population health or risk stratification

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

☐ Administrative or operational AI
☐ Other (describe below)

If other, describe:

Which entry scenario applies? (Check one.)

☐ Staff or clinician request
☐ Vendor or health record system update
☐ Leadership or board initiative
☐ Peer recommendation or network learning
☐ AI already in use (discovered during governance review)
☐ Grant-funded or externally supported

Patient populations affected:

Example: Adult primary care patients across three language groups (English, Spanish, Portuguese). Higher volume of patients with complex chronic conditions in the afternoon clinic sessions where the tool was first tested.

Health center context (size, setting, relevant details):

Example: Mid-size FQHC, urban, 14 providers across two sites.

Which governance tools did the AI Governance Hub use? (Check all that apply.)

☐ AI Signal Decision Assessment
☐ Vendor Evaluation Checklist
☐ Pre-Deployment Impact Checklist
☐ Design Plan Template
☐ AI Implementation Readiness Checklist
☐ Stewardship Review Checklist
☐ Regulatory and Compliance Monitoring Checklist
☐ Other (describe below)

If other, describe:

Key decisions made and the reasoning behind them:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Example: The AI Governance Hub approved a three-provider pilot after the AI Signal Decision Assessment showed strong alignment on Pillar 1: Systemic Signal Assessment and Pillar 3: Ecosystem Alignment. The Pre-Deployment Impact Checklist flagged that patient disclosure language had not been tested with non-English-speaking patients and the vendor had not provided accuracy data for Portuguese. The Hub approved the pilot on the condition that disclosure scripts be developed in all three languages before the first patient visit and Portuguese accuracy be reviewed at the 60-day mark.

What the AI Governance Hub did not yet know at the time of the decision:

Example: The vendor's accuracy data covered English and Spanish but grouped "other languages" together. The Hub assumed Portuguese performance would be comparable to Spanish. At 60 days, Portuguese note completeness was 12 percentage points lower than English, with health-related social needs documentation consistently omitted.

Timeline (key dates and milestones):

Example: Pilot approved January. Three-provider rollout began February. First stewardship review at 60 days (April) identified the Portuguese accuracy gap. Vendor provided a targeted update in May. Expanded to twelve providers in June after the second stewardship review confirmed the gap had closed.

What the evidence showed over time:

Example: After-hours charting dropped by approximately 40 minutes per provider per day. Note quality scores stayed consistent with pre-tool baselines in English and Spanish. Portuguese note completeness improved from 76% to 91% after the vendor update.

Unexpected outcomes:

Example: Front desk staff began fielding patient questions about the tool listening before the disclosure script was delivered, so the Hub added a waiting room notice in the first week. One provider reported the auto-generated assessment and plan section occasionally elevated a secondary diagnosis to the primary position, which was tracked as an accuracy signal and reported to the vendor.

The single most important insight from this case example:

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Example: Testing AI tools across all languages your patient population speaks, before full deployment, catches accuracy gaps that aggregate data would mask. The Portuguese documentation gap would have been invisible in a combined accuracy score.

What the AI Governance Hub would do differently:

Example: Require language-specific accuracy data from the vendor before pilot approval, rather than accepting grouped "other languages" reporting.

What other CHCs should specifically watch for:

Example: Vendor accuracy claims that group non-English languages together. Waiting room dynamics where patients learn about AI tools informally before a formal disclosure. Auto-generated clinical content that reorders diagnosis priority without provider input.

Suggested signals to monitor:

Example: Note completeness by language group (monthly). Patient questions or concerns about AI to front desk staff (ongoing). Provider override rate on auto-generated assessment and plan sections (monthly). Time-to-completion for after-hours charting (monthly, by provider).

Completing this template is internal to your health center. Sharing is optional. A health center that chooses to share emails the completed template to the Massachusetts League at MA.HCCN@massleague.org. The Massachusetts League reviews submissions for network sharing through advisories, peer learning sessions, and playbook updates.

Submitted by (role):

Date of submission:

Reviewed and approved for submission by the AI Governance Hub:

☐ Yes

☐ Pending review

> **IMPORTANT**

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Your case example strengthens the evidence base for every community health center making AI decisions. The Massachusetts League reviews all submissions and follows up if additional detail is helpful.

KEY TAKEAWAYS BY ROLE

Leadership

- Case examples build the evidence base for AI decisions across the network. Reviewing submissions before they leave the organization protects institutional learning.
- Patterns across multiple case examples inform future AI strategy. Tracking which tools succeeded and which required course correction strengthens decision-making over time.

AI Governance Hub members

- This template structures the story so other health centers can learn from it. Completing it after each governance cycle captures insights while they are fresh.
- The AI Governance Hub reviews every submission for accuracy and completeness before it goes to the Massachusetts League.

IT & Security

- Review case examples for technical accuracy before submission. Confirm that tool names, system descriptions, and integration details are correct.
- Flag patterns in tool performance data that should inform future vendor evaluations or security assessments.

WHERE THIS FITS

→ Exploring whether an AI tool is the right fit?
This template documents completed cycles. Start with Section 3: Phase 1 — Listen and Learn to capture the signal.

→ Evaluating a specific AI tool?
This template documents outcomes, not evaluations in progress. Continue with Section 4: Phase 2 — Assess and Decide.

→ Formalizing an existing AI tool?
You brought AI already in use into governance. Document what the AI Governance Hub did and what the network can learn. Complete this template.

→ Reassessing a changing AI tool?
The change triggered a fresh review. Continue with Section 4: Phase 2 — Assess and Decide, then document the outcome here when the cycle is complete.

→ Checking an approved AI tool?
You ran a stewardship review. Document the findings and share what the network can learn. Complete this template.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

→ **Replacing or exiting an AI tool?**

You managed the transition. Document what happened and what other health centers should watch for. Complete this template.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Section 7: Keeping the System Running

GOAL OF THIS SECTION

Keep governance alive after go-live: monitor the environment, keep your Playbook current, and share what you learn across the MLCHC network.

WHAT'S COVERED IN THIS SECTION

7.1 Regulatory Awareness and Proactive Monitoring
What to watch, and how often

7.2 Keeping the Playbook Current
Update your governance record

7.3 Network Feedback and Shared Learning
Give and get learning

RELATED APPENDICES

Appendix D: Tools, Templates and Resources
Companion resources for AI in healthcare

Appendix E: AI Literacy Resources
Literacy tools for staff and board

7.1 Regulatory Awareness and Proactive Monitoring

AI tools operate in a regulatory environment that keeps shifting. Federal and state policy, vendor contracts, payer requirements, and network knowledge each move on their own timelines. Watching for those changes is a standing practice your AI Governance Hub holds for as long as a tool stays active, applying the SEEDS of Innovation™ decision method it has used since Phase 1.

A full companion library for AI in healthcare, spanning federal policy, state regulation, and peer-reviewed literature, is available in Part 2: Resources for AI in Healthcare, in Appendix D: Tools, Templates and Resources.

IMPORTANT

The Regulatory and Compliance Monitoring Checklist (Phase 1 — Listen and Learn, Section 3) is where this monitoring lives. It organizes what to watch across federal, state, payer, vendor, and network sources into quarterly and monthly checks. Section 7.1 sets the accountability; the checklist carries the detail.

A change in any of these areas can trigger a fresh look at your active tools. A trigger tells you to look again, and how far back you go depends on what changed. Some changes send you through all four

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

phases; others call for a targeted re-check of a single SEEDS pillar or one phase. Your AI Governance Hub decides how deep to go, using the checklist's trigger list as its guide.

7.2 Keeping the Playbook Current

This Playbook is a living document. It reflects the decisions your health center made, the tools you deployed, and the governance practices you built. Keeping it current is how the Playbook stays useful.

What MLCHC provides.

- MLCHC reviews and updates this Playbook when new regulatory guidance, vendor developments, or network learning changes recommended practice in a meaningful way.
- MLCHC distributes updated versions to health centers through established network channels.
- MLCHC maintains version records so health centers can see what changed and when.

What your health center maintains.

- Your AI Governance Hub updates your center's local copy when your tools, policies, or governance structure change.
- Your AI Governance Hub records local decisions, any deviations from recommended practice, and the reasoning behind them.
- Your AI Governance Hub tracks which version of the Playbook your current governance rests on.

These responsibilities work together. MLCHC keeps the shared framework current; your AI Governance Hub keeps the local application accurate. Together the two layers give your health center a governance record that stays current, complete, and accountable.

7.3 Network Feedback and Shared Learning

Every health center in the MLCHC network learns alongside its peers. Each deployment decision, governance challenge, and outcome your health center experiences is relevant to peer centers facing similar choices. This subsection describes how that learning moves.

What comes to you.

- Curated case examples from peer health centers, distributed by MLCHC through established network channels.
- Updates to the shared Playbook framework, drawn from aggregated network learning.
- Regulatory and vendor alerts surfaced through the MLCHC IT Forum and network communications.

What you contribute.

- Your AI Governance Hub documents what your health center has learned, including both positive outcomes and governance challenges.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

- When MLCHC requests case examples, your AI Governance Hub responds using the Case Example Sharing Template (Phase 4 — Steward, Section 6) to structure what you share.
- Your AI Governance Hub flags emerging patterns in your signal log or monitoring activity that may be relevant to peer health centers.

Shared learning strengthens every health center in the network. What your team learns in your building becomes a resource for centers still deciding; what they share becomes guidance for your next cycle. The Case Example Sharing Template gives your contribution a consistent structure so MLCHC can gather and pass it on.

To go deeper on the resources and literacy tools that support this ongoing work, see Appendix D: Tools, Templates and Resources and Appendix E: AI Literacy Resources.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Appendix A: SEEDS of Innovation™

The decision method, in brief: the five pillars and the five levels of reasoning maturity.

This appendix is the quick reference for SEEDS of Innovation™. The Playbook's sections introduce each pillar at the point of use; come here for the five pillars side by side and the five maturity levels that show how deeply each one is answered.

What SEEDS Is

SEEDS of Innovation™ is a shared decision standard for AI in systems of care. It does more than help a health center choose better tools. **It shifts who holds the power. It moves health centers from having AI happen to them to governing it themselves, on their own terms and in service of their patients.** It works by giving everyone the same five questions to answer before they build or buy AI, so a health center can tell early whether a tool will hold up in the real world, and so the health center, not the vendor, sets the terms.

Think of SEEDS the way you already think about clinical decision-making. When a clinician sees a patient, they begin from a shared standard: a set of questions to ask, information to gather, and checks to run before reaching a diagnosis or a treatment. Standardizing those protocols, once the evidence supports them, is one of the biggest reasons healthcare quality has improved, because every patient gets the same disciplined process. SEEDS brings that same discipline to AI decisions: before a health center builds or buys AI, everyone answers the same questions first, so each decision rests on a shared standard rather than on who happens to be in the room.

SEEDS measures one thing above all: whether the real system a tool enters can carry it. A tool can be wanted, funded, and technically sound and still fail, because the system it joined had no room to hold it. SEEDS is how a health center earns that readiness on purpose, before it commits resources.

SEEDS sits between your strategy and your method. Your strategy is what you want to accomplish and why. Your method is how you put a tool in place and keep it running. SEEDS governs the decision in between: the right tool, for the right problem, in the right conditions, with the right protections, **decided by the health center, on its terms.**

Most frameworks arrive after the direction is set: they review a tool, check it against a rule, or explain why it failed once it was already in motion. **SEEDS works one step upstream, on the decision itself.** It gives every health center the same questions to answer before building or buying. When every health center asks the same questions, those questions become the standard vendors have to meet. That is how a health center shapes the field instead of inheriting it.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

The Five SEEDS Pillars

SEEDS has five pillars, answered in order. The AI Governance Hub works through them for every AI signal it evaluates. The sequence matters: each pillar builds on what the ones before it surfaced.

Pillar 1: Systemic Signal Assessment

Understand the system you are trying to change, and listen to the people closest to it, before you act.

This pillar confirms the problem being addressed is the real problem, the one the people closest to the work know is true. It reads the pattern beneath a single complaint, counts the experience of staff and patients as real evidence, and asks whether AI is even the right response, since sometimes the honest answer is a better workflow, a policy change, or more staff.

What it catches: tools built to solve the wrong problem, which let a health center do the wrong thing faster.

Pillar 2: Everyone-Centered Design

Design for the people who face the highest barriers, so the solution works for everyone.

This pillar designs for the people who face the highest barriers first: patients and staff alike, including the person who speaks a different language, the person with a disability, the person with the least time and the most reason to distrust the system. Build so the tool works for them, and it works for everyone. For AI, it confirms the tool was validated on patient populations that look like yours.

What it catches: bias hidden in the design, where a tool serves some groups well and fails others, often the people already facing the most barriers.

Pillar 3: Ecosystem Alignment

Make sure the solution fits the whole living system it joins, working outward from the person at the center.

This pillar checks that the tool holds across the whole living system, ring by ring: the person at the center, the staff who deliver care, the people who pay, the people who write the rules, and the environment that holds it all. It asks whether your technical setup can support the tool, whether your margins can absorb the cost, and whether it fits how work actually gets done.

What it catches: tools that work in the lab and fail in the real world, because conditions on the ground differ from where the tool was built.

Pillar 4: Data Sovereignty

The data belongs close to the people whose lives and work created it, along with the value it generates.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

This pillar holds two responsibilities together: your health center's responsibility to patients (their data stays within the boundaries they agreed to, in plain language, upheld over time) and your health center's responsibility to itself (what the tool uses, who has access, how data is protected, and what happens when a contract ends). Vendor data practices live here: whether a vendor trains its models on patient data, what deletion rights you hold, and what happens to data if the relationship ends.

What it catches: value taken from people who had no say, where data given by patients becomes value that flows only to companies.

Pillar 5: Stewardship

Stay responsible for the solution long after launch, and hold it to the reason you built it.

This pillar keeps a health center responsible after launch: watching what the tool does in real use, measuring it against the signal named in Pillar 1: Systemic Signal Assessment, and keeping the power to adjust, improve, or stop it. A real way to stop means a named person, a clear threshold, and a working mechanism. Stewardship guides the response to harm and the recognition of value alike: when a tool performs better than expected, that signal deserves attention and investment.

What it catches: tools launched and then left alone, drifting toward harm with no one watching.

Reasoning Maturity: How Deeply You Listen

The five questions are half of SEEDS. The other half is how deeply you answer each one. Two people can ask the same question and stand in different places, because one responds to what is loudest today and the other listens all the way down to how the system really works.

SEEDS measures this depth as reasoning maturity. It is a diagnostic rather than a scorecard: it describes how much of the system is visible to the people making the decision at the time they make it. The same question can be answered at five levels, each listening more deeply and standing on stronger evidence than the last.

Level	In one line
Level 1: Reflexive	You respond to what is visible and urgent. Your evidence is a single event happening today.
Level 2: Protective	You reduce exposure to known risks. Your evidence is the records and reviews that hold up under scrutiny.
Level 3: Normative	You follow established standards and processes. Your evidence is the official guideline everyone shares.
Level 4: Integrative	You see not just what a choice does directly, but what it triggers next, and you weigh the trade-offs between them. Your evidence is lived experience and patterns that repeat over time.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Level	In one line
Level 5: Regenerative	Each choice leaves the system stronger. Your evidence is gathered across many cycles and shared.

What Each Pillar Looks Like at Each Level

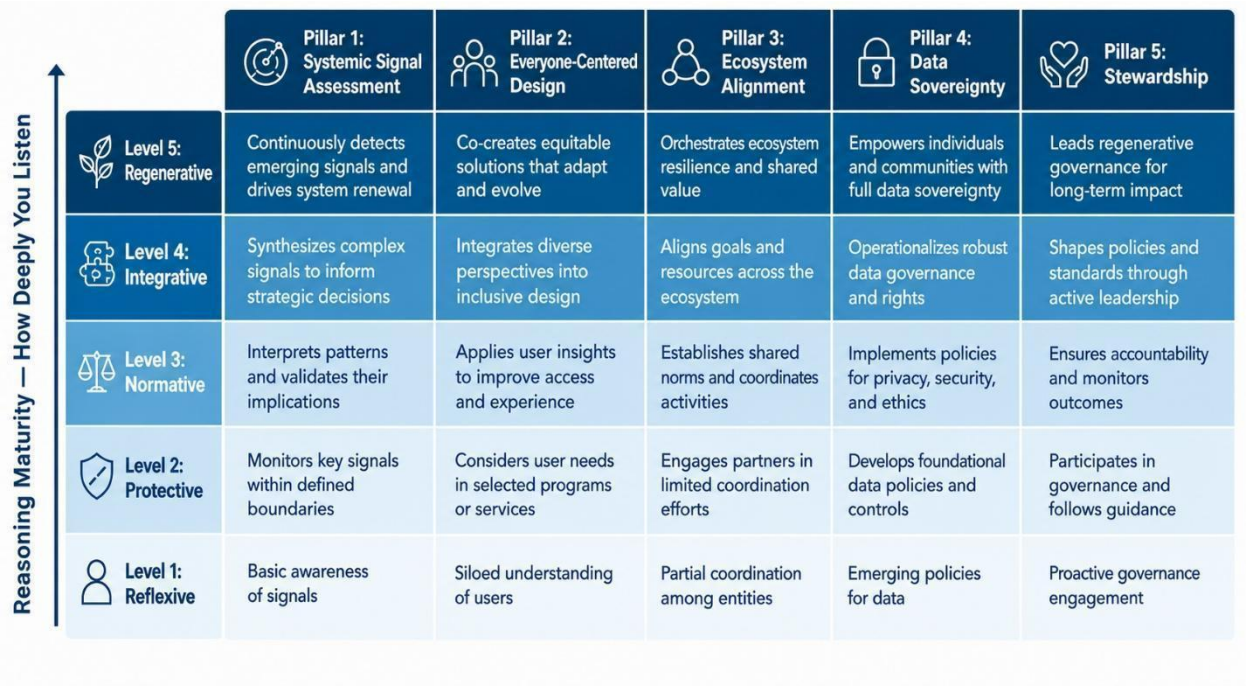

Figure A.1. SEEDS Reasoning Maturity: the five pillars across the five levels

How to Read Your Level

- Maturity lives inside each question, rather than in a separate step. You can listen deeply on one pillar and lightly on another, so a single score across all five hides more than it reveals.
- The depth of your evidence sets the size of the move you have earned. A fast read backed by one signal earns a small, reversible step; a deep read backed by real proof earns a larger commitment.
- Every health center starts at the early levels, and that is fine. When a problem is new or the evidence is thin, Level 1 or Level 2 is honest. What matters is naming what is missing: the evidence you would need to move with more confidence. That honest statement is the most valuable thing the assessment produces, and you grow into deeper levels over time.

Learn More

SEEDS of Innovation™ is developed by Phicil-itate Change. For more about the framework, visit phicil-itatechange.com.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Appendix B: Implementation Approach Reference

How This Playbook Adapted Public Frameworks for Community Health Centers

Community health centers face AI decisions shaped by frameworks built for larger, better-resourced health systems. This Playbook adapted those frameworks so they work for safety-net care.

Five tools in this Playbook draw from the public frameworks listed in this appendix. Two are Readiness tools: the Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) and the AI Implementation Readiness Checklist (Phase 3: Plan and Implement). Two are SEEDS tools that also draw from public frameworks: the Design Plan Template (Phase 3: Plan and Implement) and the Stewardship Review Checklist (Phase 4: Steward), whose Part A carries the performance and impact measures. One is a Governance tool: the Vendor Evaluation Checklist (Phase 2: Assess and Decide). Before any element from a public framework entered a Playbook tool, it was tested against eight filters designed to protect the people and teams community health centers serve.

This appendix describes the tailoring process, then presents each framework: what it is, who created it, which Playbook tools it sources, and how its content was adapted. Section 1.3 introduces this approach briefly; this appendix provides the full treatment. Additional resources for AI in healthcare appear in Appendix D.

Part 1: The CHC Tailoring Methodology

Why Tailoring Was Necessary

The public frameworks used in this Playbook were designed for health systems with dedicated AI committees, data science teams, and enterprise IT infrastructure. Most community health centers operate with lean teams, shared roles, and limited technical capacity. A framework element that protects patients in a large academic medical center may ask more of a community health center than a lean team can deliver. Governance that a team cannot sustain is governance on paper only.

The tailoring process started from a clear premise: every framework element must earn its place by serving the community health center context — protecting patients, preserving staff trust, and fitting what a lean team can actually do.

The Eight CHC Tailoring Filters

Every framework element considered for a Playbook tool was evaluated against these eight filters. A filter raises a concern when the element needs attention on that dimension. Only filters that raise a concern are documented in the evidence record.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

1. Feasibility. Can a lean community health center team do this with the expertise, time, and systems they have? If the requirement assumes specialized capacity the center lacks: shift the proof burden to the company providing the tool, simplify the requirement, or set it aside.

2. Harm and Stop-Gate Logic. Does this protect against harm to the populations community health centers serve — including patients with limited English proficiency, low literacy, LGBTQ patients, patients with sensitive diagnoses protected under 42 CFR Part 2, and others whose care requires particular attention to how AI tools are designed and used? Protective elements must be structured as hard stops with clear language.

3. Burden and Sovereignty. Does this add work to clinical or administrative teams that closes a care loop? Does it preserve the health center's control over what touches their patients, data, and protocols? Elements that generate documentation with no direct patient or staff benefit are flagged or set aside.

4. Maturity Calibration. What governance infrastructure does this assume? The Healthcare AI Governance Readiness Assessment (published by Hussein, Zink, Ramadan, and colleagues in npj Digital Medicine, 2026) defines five maturity levels for AI governance in healthcare organizations, from Level 1 (Initial/Ad Hoc) through Level 5 (Leading). Level 1 and Level 2 describe organizations with limited formal AI governance structures. Level 3 and above assume dedicated AI committees, enterprise data warehouses, and clinical informatics teams. Most community health centers operate at Level 1 or 2. Elements that assume Level 3 or 4 capacity are adapted to function at Level 1 or 2, with a clear note on what protective value is preserved or changed.

5. Scar Tissue Sensitivity. Does this resemble a pattern community health centers have struggled with: grant-funded tools that end when the grant ends, EHR rollouts that consumed staff capacity for months, compliance requirements that generated paperwork and changed nothing, or tools that failed the populations they were meant to serve? If yes, the element must name what is structurally different this time. If that question has no clear answer, the element is set aside regardless of its technical merits.

6. Evidence Legitimacy. Was this element developed for or validated in safety-net settings? If the evidence comes only from academic medical centers or large commercial systems, that gap is named plainly. The adaptation note explains how the community health center version accounts for the difference.

7. Labor and Workforce Trust. Does this respect frontline staff as decision-makers? Could it be read as monitoring, replacing clinical judgment, or reducing the role of the people doing the work? In community health centers, Federal Tort Claims Act coverage depends on providers as the final decision-makers. Any element that weakens that role is reframed or removed.

8. Collective Action Value. If every community health center applies this consistently, does it protect something the field cannot protect one center at a time? What does the field lose when centers skip this

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

step? Elements with field-level protective value are treated as essential. The cost of skipping them is paid by the individual center and by every community health center making AI decisions alongside it.

The Four Verdicts

After the filter analysis, every framework element receives one of four verdicts:

Verdict	Description
Include as-is (✓)	The element serves the community health center context and needs no changes.
Adapt (~)	The element is included with named changes. The adaptation and the reasoning behind it are documented in the Comprehensive Framework Source Register.
Simplify (↓)	The element is included at Level 1 or 2 capacity. The adaptation note records what protective value is preserved and what changes.
Exclude (✗)	The element does not belong in the community health center context. The reason is logged in the Comprehensive Framework Source Register.

The Evidence Record

Every filter analysis and verdict is documented question by question in the Comprehensive Framework Source Register — a production evidence record available to the Massachusetts League of Community Health Centers. Where this appendix synthesizes across decisions to explain the overall approach, the Register documents each individual decision with the verbatim source text, the filter analysis, the verdict, and the adaptation rationale.

Why Consistent Use Matters

This tailoring process was built so that community health centers are asking the same protective questions, applying the same standards, and holding the companies that sell AI tools to the same expectations. When one center skips a deployment gate or accepts a tool another center flagged, the gap affects the whole field. The Readiness tools in this Playbook are designed so that consistent use across health centers creates a shared standard — one that is harder for any single company or product to work around.

Part 2: Framework-by-Framework Reference

Ordered by breadth of contribution across the Playbook tools. Each entry describes the framework, identifies which Playbook tools it sources, and summarizes how the CHC Tailoring Filters shaped the adaptation.

Health Sector Coordinating Council Third-Party AI Risk Guide (April 2026)

Founding organization: Health Sector Coordinating Council, a public-private partnership between government and healthcare industry for infrastructure protection.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Source: https://healthsectorcouncil.org/ai-cyber-thirdparty/

This guide covers the full lifecycle of third-party AI risk in healthcare across seven phases, from discovery through retirement. It addresses procurement, validation, deployment, and ongoing monitoring of AI tools provided by outside companies.

Playbook tools sourced: Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) across all five pillars. AI Implementation Readiness Checklist (Phase 3: Plan and Implement) as the heaviest single source. Design Plan Template (Phase 3: Plan and Implement) as the primary source. Vendor Evaluation Checklist (Phase 2: Assess and Decide) for vendor risk evaluation. Stewardship Review Checklist (Phase 4: Steward) for ongoing governance review, including the Part A performance and impact measures.

Tailoring for community health center context: Feasibility and Maturity Calibration were the main concerns. The guide assumes a procurement and compliance infrastructure most community health centers lack. Multi-step vendor validation processes were adapted into focused checklists. Enterprise-level monitoring was simplified to lean-team review cycles. The guide's role assignment structure was adapted for shared-role environments where one person may hold multiple governance functions. Despite these adaptations, the guide's protective structure — phased deployment, required go-live conditions, incident response — was retained in full.

National Institute of Standards and Technology AI Risk Management Framework (2023)

Founding organization: National Institute of Standards and Technology, U.S. Department of Commerce.

Source: https://www.nist.gov/itl/ai-risk-management-framework

The AI Risk Management Framework organizes AI governance around four functions: Govern (establish accountability), Map (identify risks), Measure (assess risks), and Manage (respond to risks). Published in 2023, it is the most widely adopted AI governance standard in the United States.

Playbook tools sourced: Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) for risk classification and governance structure. AI Implementation Readiness Checklist (Phase 3: Plan and Implement) for governance readiness and accountability. Design Plan Template (Phase 3: Plan and Implement) for governance, management, and measurement functions. Stewardship Review Checklist (Phase 4: Steward) for governance review and, in its Part A measures, ongoing monitoring and response.

Tailoring for community health center context: Maturity Calibration and Feasibility were the main concerns. The framework assumes dedicated AI governance roles and formal risk management processes. Governance structures were adapted from standing committees to the AI Governance Hub model: a named lead with cross-functional representation and a documented decision pathway. Risk classification was simplified from a multi-tier system to a practical high, medium, and low model community health centers can apply with existing staff.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Digital Medicine Society AI Playbook

Founding organization: Digital Medicine Society (DiMe), a nonprofit professional society advancing digital medicine.

Source: https://dimesociety.org/ai-implementation-in-healthcare-playbook/

The DiMe AI Playbook provides stage-by-stage implementation guidance for AI in healthcare, covering planning, deployment, and monitoring. It includes total cost of ownership methodology and operational sustainability guidance developed for health systems adopting digital tools.

Playbook tools sourced: AI Implementation Readiness Checklist (Phase 3: Plan and Implement) as the second-heaviest source. Design Plan Template (Phase 3: Plan and Implement) for planning and deployment elements. Stewardship Review Checklist (Phase 4: Steward) for minimum monitoring standards and, in its Part A measures, implementation monitoring and total cost of ownership.

Tailoring for community health center context: Feasibility was the main concern, particularly around implementation complexity. DiMe's guidance was developed for health systems with dedicated digital health teams. Total cost of ownership methodology was retained as a protective measure: community health centers need to track the full cost of AI tools, including staff time, workflow changes, and contract obligations, to avoid adopting tools that become unsustainable. Monitoring requirements were calibrated to what lean teams can track with existing systems.

Healthcare AI Governance Readiness Assessment (Hussein, Zink, Ramadan et al., 2026)

Founding organization: Multi-institutional academic research team, published in npj Digital Medicine (2026).

Source: https://www.nature.com/articles/s41746-026-02418-7

The Healthcare AI Governance Readiness Assessment defines five maturity levels for AI governance in healthcare organizations, from Level 1 (Initial/Ad Hoc) through Level 5 (Leading). It provides the most detailed publicly available benchmark for assessing where a health system stands in AI governance capacity. This is a public-domain academic framework.

Playbook tools sourced: Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) for training readiness in Everyone-Centered Design (P2.Q2). AI Implementation Readiness Checklist (Phase 3: Plan and Implement) as the maturity calibration anchor across all domains. Stewardship Review Checklist (Phase 4: Steward) for governance maturity benchmarking.

Tailoring for community health center context: Maturity Calibration applied to nearly every element drawn from this framework. The assessment's Level 3 benchmarks describe infrastructure that exceeds what most community health centers currently have in place: dedicated AI committee, enterprise data warehouse, clinical informatics team. These benchmarks became the calibration anchor for the entire

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

tailoring process. All Readiness tool content was adjusted to function at Level 1 or 2, preserving the protective intent of higher-level requirements through simpler structures. The two worked examples in the Comprehensive Framework Source Register show this calibration in detail.

National Center for Farmworker Health Change Management Resource

Founding organization: National Center for Farmworker Health, a nonprofit organization serving community and migrant health centers.

Source: https://www.ncfh.org/

This resource provides change management guidance developed specifically for health center settings, addressing workforce transition, staff engagement, and communication during organizational change. It is one of the few change management resources built for the safety-net care context.

Playbook tools sourced: AI Implementation Readiness Checklist (Phase 3: Plan and Implement) for workforce readiness, communication, and staff engagement. Design Plan Template (Phase 3: Plan and Implement) for workforce, communication, and governance elements. Stewardship Review Checklist (Phase 4: Steward) for workforce review and, in its Part A measures, change sustainability.

Tailoring for community health center context: This content required minimal tailoring because it was developed for the health center context. Scar Tissue Sensitivity was applied as a check: confirming that the approach accounts for the specific patterns of technology adoption stress community health centers have experienced. Workforce readiness questions were adapted to address AI-specific concerns: staff uncertainty about role changes, questions about monitoring, and concerns about whether AI tools serve patients or replace staff judgment.

ONC Safety Assurance Factors for EHR Resilience (SAFER) Guides

Founding organization: Office of the National Coordinator for Health Information Technology, U.S. Department of Health and Human Services.

Source: https://www.healthit.gov/topic/safety/safer-guides

The Safety Assurance Factors for EHR Resilience guides drawn on here are Computerized Provider Order Entry with Decision Support, System Management, and Organizational Responsibilities. Originally designed for electronic health record safety, they extend to AI tools that interact with clinical systems and workflows.

Playbook tools sourced: Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) for clinical safety questions, including alert fatigue, override protocols, system labeling, and patient information in multiple languages. Stewardship Review Checklist (Phase 4: Steward) for ongoing safety monitoring.

Tailoring for community health center context: Feasibility was the main concern around staffing and process assumptions. The guides assume health IT departments with formal change management

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

processes. Clinical safety protections were kept in full — these are patient-facing protections community health centers must maintain regardless of team size. Implementation requirements were adapted from enterprise IT workflows to lean-team protocols. System management practices were reframed for settings where one or two people manage the full technology stack.

ECRI (Incorporating AI into Healthcare, June 2024; Ethical Use of AI in Healthcare, August 2024; Managing the Risks of AI in Healthcare, February 2025)

Founding organization: ECRI (formerly the Emergency Care Research Institute), an independent nonprofit patient safety organization.

Source: https://home.ecri.org/pages/ai-resources

Three complementary documents address AI safety from different angles: Incorporating AI into Healthcare (organizational conditions for safe deployment, including total-systems-safety principles), Managing the Risks of AI in Healthcare (identifying and reducing the risks of AI-enabled tools), and Ethical Use of AI in Healthcare (governance and accountability).

Playbook tools sourced: Pre-Deployment Impact Checklist (Phase 2: Assess and Decide) for rollback protocols and incident response requirements. Stewardship Review Checklist (Phase 4: Steward) for safety review and, in its Part A measures, safety monitoring and incident tracking.

Tailoring for community health center context: Feasibility was the main concern around assumed staffing and reporting capacity. ECRI's safety standards were developed for health systems with dedicated patient safety officers and formal incident reporting. Patient safety requirements were kept as required elements — adapted for lean teams with fewer formal roles and simpler reporting, while preserving the same conditions that must be met before deployment. Total systems safety principles were built into the STOPGATE structure of the Pre-Deployment Impact Checklist.

Donabedian Quality Model

Founding organization: Developed by Avedis Donabedian (University of Michigan School of Public Health), widely adopted in healthcare quality measurement since 1966.

Source: *Academic model; see Appendix D for additional resources.*

The Donabedian model evaluates healthcare quality across three domains: Structure (what resources and systems are in place), Process (how care is delivered), and Outcome (what results for patients). Applied in this Playbook as the measurement framework for AI performance.

Playbook tools sourced: Design Plan Template (Phase 3: Plan and Implement) for the five-pillar Evidence Plan, extending Donabedian's original three categories to include workforce, financial, and governance dimensions. Stewardship Review Checklist (Phase 4: Steward) as the primary measurement framework for the five accountability questions in its Part A measures, and for performance-domain review.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Tailoring for community health center context: The model required minimal structural change because it is a way of thinking about measurement, rather than a set of implementation steps. Feasibility shaped the choice of starter measures within each category: community health centers select from measures they already collect or can add with existing systems. The extension from three categories to the five SEEDS pillars ensures the framework captures workforce burden, financial sustainability, and governance functioning — dimensions especially relevant to how community health centers experience AI tools.

Institute for Healthcare Improvement Plan-Do-Study-Act Cycle

Founding organization: Institute for Healthcare Improvement, a nonprofit focused on healthcare quality improvement.

Source: https://www.ihi.org/library/tools/plan-do-study-act-pdsa-worksheet

The Plan-Do-Study-Act cycle is a structured approach to testing and refining changes in healthcare settings, widely used in community health centers for quality improvement work.

Playbook tools sourced: Design Plan Template (Phase 3: Plan and Implement) as the structural foundation of the Pilot Charter and Living Record sections. Stewardship Review Checklist (Phase 4: Steward) for the review-cycle structure and the continue, adjust, or stop decision.

Tailoring for community health center context: The Plan-Do-Study-Act cycle required minimal adaptation because community health centers already use it for quality improvement. Applying it to AI implementation connects AI governance to a process teams already know. The Pilot Charter structures a prediction-based test: what do we expect to happen, what actually happened, what do we change. The Living Record carries that learning forward through successive review rounds. This familiar structure reduces the learning curve of AI governance for teams who have used Plan-Do-Study-Act in other improvement work.

Coalition for Health AI Responsible AI Guide (2024)

Founding organization: Coalition for Health AI, a multi-stakeholder initiative convening health systems, technology companies, and researchers to establish standards for responsible health AI.

Source: chai.org/workgroup/responsible-ai/responsible-ai-guide-raig-and-raig-executive-summary

The Responsible AI Guide serves as a playbook for the development and deployment of AI in healthcare, providing actionable guidance on ethics and quality assurance across the AI lifecycle.

Playbook tools sourced: Stewardship Review Checklist (Phase 4: Steward) for governance and adverse-event review, and for governance and monitoring principles in its Part A measures.

Tailoring for community health center context: Evidence Legitimacy and Feasibility were the main concerns. The guide was developed for health systems with the resources to evaluate AI tools at a

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

research level. Its governance and quality-assurance standards were translated into operational criteria community health centers can apply with existing staff, and rewritten from research language into plain language a lean team can act on during a review cycle.

Supporting Frameworks

These frameworks contribute to specific domains within the Playbook tools. They reinforce or extend the primary frameworks listed above.

National Association of Community Health Centers Action Guides and Value Transformation Framework

Founding organization: National Association of Community Health Centers, the national membership organization for federally funded health centers.

Source: https://www.nachc.org/about-nachc/our-work/quality-center/elevate/

The Action Guides provide strategic guidance for health center quality improvement. The Value Transformation Framework connects improvement work to the broader shift toward value-based care. Together, they offer a community-health-center-specific lens on how to measure whether changes are making a difference for patients and communities.

Playbook tools sourced: Stewardship Review Checklist (Phase 4: Steward), Part A, for patient engagement measures, outcome measurement, and alignment with the value goals community health centers already pursue.

Tailoring for community health center context: This content was developed for community health centers and required minimal adaptation. It reinforced the requirement that AI performance measurement connect to patient-reported outcomes and community benefit — goals health centers are already accountable for through their federal funding.

Betsy Lehman Center for Patient Safety — Including the Patient Voice: A Guide to Engaging the Public in Programs and Policy Development (Revised 2024)

Founding organization: Betsy Lehman Center for Patient Safety, a Massachusetts state agency focused on patient safety improvement.

Source: https://betsylehmancenterma.gov/tools/including-the-patient-voice

This guide provides practical approaches for engaging patients and community members in program and policy decisions. The Playbook applies its public-engagement methods to patient involvement in AI governance.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Playbook tools sourced: AI Implementation Readiness Checklist (Phase 3: Plan and Implement) for the patient advisory board briefing requirement. Design Plan Template (Phase 3: Plan and Implement) for patient experience criteria in the Pilot Charter.

Tailoring for community health center context: This guidance aligns with a strength community health centers already have: deeper community relationships than most health systems. The patient engagement approach was adapted to build on existing patient advisory structures rather than creating new ones. Community health centers that already have functioning patient boards or advisory committees can extend those relationships to cover AI governance questions.

World Health Organization Ethics and Governance of Artificial Intelligence for Health (2021)

Founding organization: World Health Organization.

Source: https://www.who.int/publications/i/item/9789240029200

This guidance establishes six ethical principles for health AI: protect autonomy, promote well-being, ensure transparency, foster accountability, ensure inclusiveness, and promote responsiveness and sustainability. It provides a global framework for ethical AI use in healthcare settings.

Playbook tools sourced: Stewardship Review Checklist (Phase 4: Steward) for patient-centered review and, in its Part A measures, patient experience and outcome consistency across the populations a health center serves.

Tailoring for community health center context: The ethical principles were translated into concrete questions a health center team can answer during a review cycle, rather than aspirational statements. The principle of inclusiveness was operationalized through the requirement that AI performance be reviewed across patient populations served by the center — a practical check that connects global ethics to local accountability.

Kaplan and Norton Balanced Scorecard

Founding organization: Developed by Robert Kaplan and David Norton (Harvard Business School), widely adopted in strategic performance management since 1992.

Source: *Academic model; see Appendix D for additional resources.*

The Balanced Scorecard measures organizational performance across four perspectives: financial, customer, internal process, and learning and growth. It ensures that strategic decisions are evaluated across multiple dimensions rather than a single metric.

Playbook tools sourced: Stewardship Review Checklist (Phase 4: Steward) for mission alignment and cost review, and, in its Part A measures, strategic alignment, sustainability measurement, and cost-benefit analysis.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Tailoring for community health center context: The corporate framing was adapted for mission-driven organizations. 'Customer perspective' became patient and community experience. 'Financial perspective' became sustainability and total cost of ownership. The four perspectives ensured the Stewardship Review Checklist's Part A measures assess AI performance across the dimensions health center leadership already tracks, rather than introducing a separate measurement system.

Kotter Eight-Step Change Model

Founding organization: Developed by John Kotter (Harvard Business School), a widely used model for leading organizational change since 1996.

Source: *Academic model; see Appendix D for additional resources.*

The model provides a sequential approach for leading change, from building a sense of purpose through anchoring new practices into how the organization operates day to day.

Playbook tools sourced: Stewardship Review Checklist (Phase 4: Steward) for ongoing governance review and, in its Part A measures, sustainability and change anchoring.

Tailoring for community health center context: Kotter's model assumes dedicated change management leadership. In community health centers, change leadership is distributed across existing roles. The model's anchoring principles were built into the sustainability and governance review questions, so that AI governance becomes part of ongoing operations rather than a standalone initiative that requires its own leadership structure.

HRSA Section 330 Requirements

Founding organization: Health Resources and Services Administration, U.S. Department of Health and Human Services.

Source: https://bphc.hrsa.gov/compliance/compliance-manual/introduction

Section 330 of the Public Health Service Act establishes the federal requirements governing Federally Qualified Health Centers, including governance structure, patient participation, and service delivery standards.

Playbook tools sourced: AI Implementation Readiness Checklist (Phase 3: Plan and Implement) for patient board and advisory committee briefing. Design Plan Template (Phase 3: Plan and Implement) for patient engagement and community governance elements in the Evidence Plan.

Tailoring for community health center context: Section 330 is the regulatory foundation community health centers already operate under. Its patient participation requirements were connected to AI governance: if a health center is putting an AI tool in place that touches patient care, the patient board or advisory committee should be briefed. This is a compliance connection that links existing federal

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

requirements to new AI governance activities, rather than adding a separate layer of patient engagement.

Complementary Framework

The National Academy of Medicine has published guidance on AI ethics and governance in healthcare settings. This work informed the principled design of the Playbook's approach to AI governance and is cited in Appendix D. The Readiness tools draw their operational content from the primary and supporting frameworks listed above; the National Academy of Medicine serves as a complementary reference for the broader policy context in which community health center AI governance operates.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Appendix C: Glossary

Key Terms and Acronyms

This glossary is a reference for terms and acronyms used throughout the Playbook, as well as in its tools and appendices. Entries are alphabetical. When relevant, links to external organizations are hyperlinked to provide additional detail and context. External links can change, so hyperlinks are current as of completion of this document.

ACRONYM	TERM	DEFINITION
	42 CFR Part 2	Federal regulation on the confidentiality of records for patients receiving substance use disorder (SUD) treatment. Applies to health centers that provide or coordinate SUD services. Requires specific patient authorization before records can be shared, with stricter controls than standard HIPAA requirements.
	AI Governance Hub	The standing body inside your health center that reviews AI signals, makes decisions, tracks outcomes, and holds the governance process together over time. The AI Governance Hub is not a one-time committee — it is an ongoing governance function. The AI Governance Hub applies the five SEEDS pillars to every AI signal it reviews.
	AI literacy	Understanding enough about AI to ask the right questions, recognize AI when you encounter it, and evaluate whether to trust what it produces. Three levels cover most needs: awareness, navigation, and decision-making.
	AI Signal Decision Assessment	The main tool the AI Governance Hub uses to decide whether to move forward on an AI signal. It walks through all five SEEDS pillars in order, asking a set of questions for each one, and ends with a clear recommendation: Approve, Approve with conditions, Hold, or Do not proceed.
BAA	Business Associate Agreement	A legally required contract with any vendor or partner that handles protected health information (PHI) on your health center's behalf. The BAA defines what data the vendor can access, how it must be protected, and what happens to the data when the relationship ends. AI vendors that access, process, or store PHI require a BAA before any data sharing begins.
BYOD	Bring Your Own Device	A policy that allows staff to use personal devices — phones, tablets, or laptops — to access the health center's networks, applications, or data. BYOD creates governance considerations in

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
		AI because personal devices may enable staff to use AI tools that have not been reviewed or approved by the health center.
CHAI	**Coalition for Health AI**	A multi-stakeholder organization focused on guidance and standards for responsible AI in healthcare. CHAI's core principles for trustworthy AI include usefulness and effectiveness, freedom from bias, safety, transparency and accountability, and privacy and security. CHAI publishes the Blueprint for Trustworthy AI, the Responsible AI Guide (RAIG), and the Responsible AI Implementation Checklist (RAIC), all freely available.
CHC	**Community Health Center**	A nonprofit clinic located in a medically underserved area that provides comprehensive primary, preventive, and dental care — along with behavioral health and substance use services — to everyone, regardless of insurance status or ability to pay. This Playbook was built specifically for community health centers.
CMS	**Centers for Medicare and Medicaid Services**	The Centers for Medicare and Medicaid Services (CMS) is a federal agency that provides health coverage through Medicare, Medicaid, the Children's Health Insurance Program, and the Health Insurance Marketplace.
	Cognitive burden	The operational risk that an AI tool adds mental load to a workflow without removing it — requiring more steps, more alerts, or more decisions from staff who are already working at full capacity. Cognitive burden is an Ecosystem Alignment (SEEDS Pillar 3) risk. A tool that saves time on one task but creates burden elsewhere has not produced clean value. Cognitive burden is distinct from cognitive surrender.
	Cognitive surrender	The tendency to accept AI-generated output without applying independent judgment, because the AI answered with confidence, speed, or apparent authority. Cognitive surrender is addressed by the maturity check within the AI Signal Decision Assessment and is revisited over time through stewardship. Cognitive surrender is distinct from cognitive burden.
CPG	**Commonwealth Purchasing Group**	The Commonwealth Purchasing Group (CPG) was founded by MLCHC. CPG is a group purchasing organization established to address economic issues facing health centers and nonprofit community service organizations. CPG provides collective negotiation as leverage when evaluating vendor terms and pricing.
	Data sovereignty	A patient's right to understand how their data is used, who can access it, and what happens to it as it moves through AI-enabled tools and systems. In the context of the SEEDS of Innovation™ framework, Data Sovereignty is the fourth pillar. It applies this principle to every AI tool decision: who owns the data, what

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
		secondary uses are permitted, what happens to patient data if the vendor relationship ends, and what rights patients have to request deletion.
DiMe	**Digital Medicine Society**	A multi-stakeholder organization that publishes AI implementation guidance, including a practical deployment guide with a dedicated low-resource settings companion track. DiMe's maturity framework for health AI includes five levels and an evaluation readiness self-assessment. Freely available.
DUA	**Data Use Agreement**	A legal contract that governs how data from a specific source can be used, shared, and protected. In the context of AI, a DUA may be required when a health center shares patient data — even in de-identified form — with a research partner, network, or vendor.
	Division of Insurance	Within the Massachusetts State Office of Consumer Affairs and Business Regulation, the Division of Insurance monitors matters related to insurance through several departments. The Division protects consumer interests by providing accurate and unbiased information so consumers may make informed decisions and by intervening on behalf of consumers who believe they have been victimized by unfair business practices.
	Ecosystem Alignment	The third pillar of SEEDS of Innovation™. The Ecosystem Alignment pillar ensures that the solution fits and holds across the whole living system it joins, from the person at the center out to the environment itself.
EHR	**Electronic Health Record**	A secure digital system that stores and manages patient health information, including clinical notes, diagnoses, medications, and test results. Many AI tools are embedded in or integrated with EHR systems, meaning an AI feature may activate as part of a routine platform update without a separate procurement decision.
EOHHS	**Executive Office of Health and Human Services**	This executive office is comprised of 11 agencies and the MassHealth program for the state of Massachusetts. The office provides access to medical and behavioral health care, substance use disorder treatment, long term services and supports, as well as nutritional and financial benefits.
	Everyone-Centered Design	The second pillar of SEEDS of Innovation™. The Everyone-Centered Design pillar asks whether the people affected by an AI tool — patients, staff, and the communities the health center serves — were included in the decision and whether the tool works for them. It is the governance mechanism that ensures patient and staff voice reaches every AI decision the health center makes.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
FQHC	**Federally Qualified Health Center**	A community health center that has received federal certification under the Social Security Act and meets specific requirements set by the Health Resources and Services Administration (HRSA). FQHCs receive enhanced Medicare and Medicaid reimbursement and operate under a consumer-majority board governance model.
FedRAMP	**Federal Risk and Authorization Management Program**	A government-wide program that provides a standardized approach to security and risk assessment for cloud products and services.
	G.L. c. 93H	Massachusetts General Law Chapter 93H. A Massachusetts state law governing the responsibilities of organizations that experience a data security breach involving personal information. Requires notification of affected individuals, the state Attorney General, and the Office of Consumer Affairs and Business Regulation when a breach occurs.
	Go-live	The day an AI tool officially starts being used in real work, with real patients or staff.
HCCN	**Health Center Controlled Network**	A network organization that provides health IT support, training, and technical assistance to community health centers. HCCNs are funded by HRSA and often serve as the channel through which health centers access technology infrastructure, data resources, and peer learning. In AI governance, the HCCN may serve as a gatekeeper for technology procurement and shared systems decisions.
HEDIS	**Healthcare Effectiveness Data and Information Set**	A set of standardized measures used for performance improvement, accountability, and benchmarking in healthcare. HEDIS metrics are managed by the National Committee for Quality Assurance (NCQA). Some AI tools are designed to support HEDIS measure performance — understanding what a tool measures matters for evaluating its fit with your health center's goals.
HIPAA	**Health Insurance Portability and Accountability Act**	Federal legislation that protects patient health information privacy and governs how protected health information (PHI) can be used, stored, and shared. All AI tools that interact with PHI must comply with HIPAA, and vendors must sign a Business Associate Agreement (BAA) before handling PHI.
HITECH	**Health Information Technology for Economic and Clinical Health Act**	The HITECH Act promotes the adoption and meaningful use of health information technology. It focuses on privacy and security concerns related to electronic transmission of health information, as an enhancement to HIPAA regulations.
HITRUST	**Health Information Trust Alliance**	The HITRUST certification is a framework to assess cybersecurity and risk management. This framework helps organizations to ensure compliance with various data protection requirements.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
HHS	**Health and Human Services**	HHS is a federal department that is responsible for maintaining health protections, research, food safety, health care, aging, and more.
HRSA	**HRSA Section 330**	The Health Resources and Services Administration (HRSA) is within the U.S. Department of Health and Human Services (HHS). Section 330 of the Public Health Service Act establishes the federal requirements governing Federally Qualified Health Centers, including governance structure, patient participation, and service delivery standards.
	Massachusetts 201 CMR 17.00	A Massachusetts state regulation that sets minimum standards for the protection of personal information held by organizations doing business with Massachusetts residents. Requires written information security programs and specific safeguards for data at rest and in transit.
	MassHealth QEIP	MassHealth Quality and Equity Incentive Program. A Massachusetts state-level program that includes a restriction on the use of AI for medical interpretation within the MassHealth system. Health centers serving patients through MassHealth should review current QEIP requirements before deploying AI in workflows that involve interpretation or medical decision support.
	Minimum Viable Governance	The five components every health center needs in place before approving any AI tool: an AI tool inventory, a named decision authority, an AI policy, a staff disclosure mechanism, and a defined review cadence. Minimum Viable Governance is the floor, not the ceiling — the structure that makes the next AI decision governable. The Playbook does not abbreviate this term — write it out.
MLCHC	**Massachusetts League of Community Health Centers**	Also referred to as the League. The statewide primary care association for community health centers across Massachusetts. MLCHC's AI strategy works at the network level — shared learning, shared standards, and collective action. This Playbook was developed through an engagement with MLCHC to bring that strategy into health center operations.
NACHC	**National Association of Community Health Centers**	The national membership and advocacy organization for community health centers. NACHC publishes the AI Action Guide, a plain-language orientation to AI for CHC leaders and staff, and formally endorses the NIST AI Risk Management Framework and CHAI guidance as standards for CHC AI governance. Freely available through NACHC.
NIST	**National Institute of Standards and Technology**	A federal agency that develops standards and frameworks. NIST's AI Risk Management Framework (AI RMF) is the most widely

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
		referenced AI governance framework across healthcare, organized around four functions: Govern, Map, Measure, and Manage. NACHC endorses it as a foundational standard for CHC AI governance. Freely available.
ONC	**Office of the National Coordinator for Health Information Technology**	A federal office within the U.S. Department of Health and Human Services (HHS). ONC leads HHS' efforts on health IT and is a resource to the entire health system to support the adoption of health IT. ONC also promotes national standards-based health information exchange to improve healthcare.
PHI	**Protected Health Information**	Any health information that can identify an individual and relates to their past, present, or future health condition, healthcare services, or payment for healthcare. Under HIPAA, PHI must be protected whether in electronic, paper, or spoken form. AI tools that access, generate, or process patient information are likely handling PHI and require appropriate safeguards.
PII	**Personally Identifiable Information**	Information that can be used to identify a specific individual, either alone or in combination with other data. PII is a broader category than PHI and includes names, addresses, social security numbers, and similar identifiers. Some state laws — including Massachusetts 201 CMR 17.00 — regulate the protection of PII separately from federal HIPAA requirements.
	Pilot charter	A short document required before any limited test of an AI tool begins. The pilot charter names the purpose, scope, population affected, owner, timeline, guardrails, data use expectations, training expectations, and measures. The charter prevents an informal exploration from becoming a de facto adoption.
	Reasoning maturity	Maturity describes the strength of the evidence a decision rests on — how much of the system is visible to the people making the decision at the time they make it. It is a diagnostic measurement, not a scorecard. Reasoning maturity is a critical aspect of SEEDS.
RAIC	**Responsible AI Implementation Checklist**	A resource published by CHAI designed to guide organizations in evaluating an AI tool or system against CHAI's standards for trustworthy AI. Freely available.
RAIG	**Responsible AI Guide**	A resource published by CHAI that incorporates practices for the development and deployment of AI tools in healthcare. Freely available.
RCM	**Revenue Cycle Management**	The financial process of managing patient service revenue from the time care is scheduled or requested through final payment. AI tools are increasingly used in revenue cycle workflows — including coding, billing, prior authorization, and claim management. Revenue cycle AI decisions carry compliance, accuracy, and

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
		workforce implications and should go through the same governance process as clinical AI.
ROI	**Return on Investment**	A measure used to assess whether an investment produced meaningful value relative to its cost. In AI governance, a complete picture of value asks not just whether a tool saves time or reduces cost, but whether it is working for patients and staff, whether it creates burden elsewhere in the workflow, and whether it is improving outcomes over time.
	Section 1557 of the Affordable Care Act	Section 1557 is a provision within the Affordable Care Act which prohibits discrimination on the grounds of race, color, national origin, sex, age, or disability in certain health programs and activities. Section 1557 applies to any health program, organization or activity that receives funding from HHS.
SEEDS	**SEEDS of Innovation™**	A decision framework developed by Phicil-itate Change that structures how a health center evaluates any AI signal. SEEDS consists of five pillars applied in sequence: Systemic Signal Assessment, Everyone-Centered Design, Ecosystem Alignment, Data Sovereignty, and Stewardship. Each pillar surfaces a different dimension of the decision. SEEDS is the decision method running through this Playbook.
	Shadow AI	AI tools that are in active use inside the health center but have never been reviewed or approved through a formal governance process. Shadow AI often arrives through platforms the health center already trusts — a feature activated inside an existing tool, or a commercially available AI application a staff member found useful. The governance response is to bring shadow AI through the same review process as any other signal.
	Signal	Any indication that something AI-related needs your AI Governance Hub's attention. Sources include vendor outreach, staff observations, patient feedback, regulatory or payer changes, vendor or health record system feature updates, peer health center experiences, and AI that arrived without a governance decision. Signals are not emergencies — they are information that triggers the Playbook's five-phase process.
SLA	**Service-Level Agreement**	A written agreement between a health center and a vendor that defines the standard of service to be provided — uptime, response times, support availability, and what happens when the vendor fails to meet those standards. For AI tools, the SLA should also address what happens during model updates, retraining events, or performance degradation.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

ACRONYM	TERM	DEFINITION
	Stewardship	The fifth pillar of SEEDS of Innovation™. The Stewardship pillar names what comes after the decision: who is accountable for the tool over time, how performance gets measured, and what conditions trigger a pause, redesign, or stop. Stewardship treats the decision to approve as the beginning of the governance work, not the end.
	STOPGATE	A binding condition that halts the go-live process until it is resolved, because going forward would create harm or exposure that cannot be undone after the fact. In the Stewardship Review Checklist, a Yes on a policy and boundary item escalates to the AI Governance Hub before the tool continues in use.
	SOC 2 Type II certification	SOC 2 assessment defines criteria for managing customer data using five trust service principles: security, availability, processing integrity, confidentiality and privacy. SOC 2 Type II certification provides a determination of the effectiveness of an organization's security controls over an extended period of time on how effectively the systems operate to keep customer data safe.
	Systemic Signal Assessment	The first pillar of SEEDS of Innovation™. The Systemic Signal Assessment helps to understand the system you are trying to change, and to listen to the people closest to it, before acting. Before you choose what to build or buy, you study the system you are trying to change, and what it is telling you.
UDS	**Uniform Data System**	A reporting system used by HRSA-funded FQHCs to document patient demographics, services provided, clinical processes, and health outcomes. UDS data measures performance and accountability across the federally qualified health center program. Some AI tools are designed to support UDS reporting or to flag care gaps tied to UDS measures.
	Zombie pilots	AI tools left running, spreading, or abandoned without a scheduled governance decision: pilots that become permanent, keep expanding, or fade away by habit instead of by decision — leaving behind incomplete records, staff with inconsistent practices, and ongoing data access that was never formally closed. A formal exit and decommissioning process is designed to prevent zombie pilots.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Appendix D: Tools, Templates and Resources

Index of tools, templates and external resources

Part 1: Playbook Tool and Template Index

This index lists all fourteen tools and templates referenced throughout the Playbook. Tools and templates are listed in order of appearance by section.

1	Approved-Use Inventory Checklist	How do we build and maintain our AI tool inventory?	Section 2
2	AI Policy Checklist	What does our AI policy need to cover?	Section 2
3	AI Governance Hub Getting Started Kit	How do we set up AI governance?	Section 2
4	Leadership Talking Points Template	What do we say about this AI decision, and who needs to hear it?	Section 2
5	AI Signal Intake Form Template	How do we capture and route an AI signal?	Section 3
6	Patient AI Feedback Form Template	How do we hear from patients about AI in their care?	Section 3
7	Regulatory and Compliance Monitoring Checklist	What do we watch for, and how often?	Section 3
8	AI Signal Decision Assessment	Should we move forward, and how strong is our reasoning?	Section 4
9	Vendor Evaluation Checklist	What do we need to know before and during vendor evaluation?	Section 4
10	Pre-Deployment Impact Checklist	Is this AI tool safe to deploy?	Section 4
11	Design Plan Template	How do we design a plan for implementation that works for our patients, our staff, and our health center?	Section 5
12	AI Implementation Readiness Checklist	Is our organization ready to launch this?	Section 5
13	Stewardship Review Checklist	Is this tool still delivering what we said it would — what do the measures show, and what do we do about it?	Section 6
14	Case Example Sharing Template	How do we share what we learned?	Section 6

Part 2: Resources for AI in Healthcare

This section provides some external resources that are available to help stay current on AI in healthcare, particularly for community health centers (CHC). These resources can be referenced as CHCs build, govern, and advance their healthcare-related AI programs. Resources are organized into six categories: AI Governance, AI Frameworks, AI Use Cases, AI Trends, AI Regulation and Policy, and Relevant Literature. For AI literacy training courses and curricula, see Appendix E: AI Literacy Resources.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

This is not an extensive list and should be regularly reviewed to keep pace with rapid change and evolution in the field of AI. Links are verified as of publishing (June 2026). Your Health Center Controlled Networks (HCCN) or AI Governance Hub members can be a resource with questions.

AI Governance

Foundational frameworks and principles that guide how health centers make, oversee, and document decisions about AI tools — from initial vendor evaluation through ongoing monitoring.

Resource and Organization	Audience	Description and Direct Link
AI Risk Management Framework (AI RMF 1.0) *National Institute of Standards and Technology (NIST)*	All roles	The foundational U.S. framework for managing AI risks across the full lifecycle. Covers governance, mapping, measuring, and managing AI risks. Widely adopted in health care settings. → airc.nist.gov
Blueprint for an AI Bill of Rights *White House Office of Science and Technology Policy (OSTP)*	Leadership · Board	Five principles guiding responsible AI use: safe systems, non-discriminatory tools, data privacy, notice and explanation, and human alternatives. Directly applicable to patient-facing AI tools. → whitehouse.gov/ostp/ai-bill-of-rights
Assurance Standards Guide for Healthcare AI (v1.0) *Coalition for Health AI (CHAI)*	AI leads · Clinical leads	Practical standards for evaluating, validating, and monitoring AI tools in health care. Developed for health systems and safety-net providers. Includes implementation guidance aligned with NIST AI RMF. → coalitionforhealthai.org/resources
Ethics and Governance of Artificial Intelligence for Health *World Health Organization (WHO)*	Leadership · Policy roles	WHO's six core principles for ethical AI in health: autonomy, well-being, sustainability, accountability, transparency, and equity. Useful for grounding governance policy in internationally recognized standards. → who.int — AI for Health guidance
AI Governance Principles for Health Care *American Hospital Association (AHA)*	Leadership · Board	AHA's guidance on governance structures, vendor accountability, data stewardship, and community-benefit considerations for hospitals and health systems deploying AI. → aha.org — AI Governance Principles
AI Policy and Principles *American Medical Association (AMA)*	Clinical leadership	AMA's formal policies on physician oversight of AI, algorithmic transparency, liability frameworks, and patient rights. Practical reference when developing provider-facing AI governance policies. → ama-assn.org — AI Principles

AI Frameworks

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Structured approaches to decision-making, change management, and organizational readiness as health centers move from exploring AI to launching and scaling it.

Resource and Organization	Audience	Description and Direct Link
Responsible AI in Health Care: A Framework for Boards *Institute for Healthcare Improvement (IHI)*	Board · Leadership	Structured guidance for governing bodies on overseeing AI adoption, asking the right questions, and building accountability into AI decisions from the top. → ihi.org — AI Governance for Boards
Health AI Landscape Report *National Association of Community Health Centers (NACHC)*	All roles	CHC-specific analysis of AI readiness, current adoption patterns, and recommended priority areas for federally qualified health centers. The closest peer benchmark available for MLCHC member organizations. → nachc.org — AI Resources
AI Change Management Toolkit *Agency for Healthcare Research and Quality (AHRQ)*	Operations · QI leads	Practical tools for introducing AI into care workflows, including readiness assessments, staff engagement guides, and change resistance frameworks designed for clinical settings. → ahrq.gov — AI Tools & Resources
AI Adoption Framework for Health Systems *McKinsey Center for Societal Benefit through Healthcare*	Leadership · Strategy	A structured approach to moving from AI experimentation to scaled adoption, with attention to workforce impact, governance, and return-on-investment considerations in resource-constrained environments. → mckinsey.com — AI in Health Care
Transforming Health Care with AI: The Impact on the Workforce and Organizations *Harvard Business Review (HBR)*	Leadership	Research-based perspective on how AI restructures clinical and administrative roles, with guidance on workforce planning, transparency, and organizational culture during AI transitions. → hbr.org — AI and Machine Learning
AI in Primary Care: A Playbook *Primary Care Collaborative (PCC)*	Clinical leadership · QI leads	Implementation guidance specific to primary care settings — scheduling optimization, care gap identification, documentation assistance, and patient engagement tools — with access and affordability considerations built in. → pcpcc.org — AI in Primary Care

AI Use Cases

Evidence-based examples of AI applications in clinical care, operations, behavioral health, population health, and patient engagement. These use cases have a particular application to CHC settings.

Resource and Organization	Audience	Description and Direct Link

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI and Health Equity: Use Cases for Safety-Net Providers *National Health Law Program (NHeLP)*	All roles	Examines AI use cases through an equity lens, including risk stratification, population health, and screening tools — with particular attention to bias and access concerns for underserved populations. → healthlaw.org — AI and Health Equity
AI Use Case Compendium *Office of the National Coordinator for Health IT (ONC)*	Operations · IT leads	Catalogues real-world AI applications across ambulatory, hospital, and community health settings — organized by domain (clinical decision support, administrative efficiency, patient engagement, population health). → healthit.gov — AI Initiatives
NEJM AI — Clinical Research and Case Studies *New England Journal of Medicine (NEJM AI)*	Clinical leadership · Quality	Peer-reviewed AI clinical research covering diagnostics, chronic disease management, behavioral health screening, and care coordination. The most rigorous source for evaluating clinical evidence behind AI tools. → ai.nejm.org
AI in Federally Qualified Health Centers: Innovation Spotlights *Health Resources and Services Administration (HRSA)*	All roles	HRSA's collection of innovation case studies from health center grantees, including AI-assisted care management, telehealth, and EHR optimization projects. Searchable by population served and technology type. → hrsa.gov — Innovation Spotlights
Rock Health Annual Digital Health Consumer Survey *Rock Health*	Leadership · Strategy	Annual survey of consumer attitudes toward digital health and AI, broken out by race, income, insurance status, and geography — critical for understanding patient perspectives in MLCHC member communities. → rockhealth.com/research
AI for Behavioral Health: Resources and Tools *Substance Abuse and Mental Health Services Administration (SAMHSA)*	Behavioral health leads	Behavioral health-specific AI resources covering crisis response tools, risk stratification, documentation assistance, and 42 CFR Part 2 compliance considerations for AI tools handling substance use records. → samhsa.gov — Digital Health
Digital Innovations at Community Health Centers: AI Ambient Scribing in Rural Health Care *The Commonwealth Fund*	Clinical leadership	Case study exploring AI scribing and implications for community health centers (CHC). Describes the AI digital divide that is emerging in health systems and CHCs that provide critical access to care. → Case study — AI scribing at CHCs

AI Trends

Current intelligence on where healthcare AI is headed; including investment trends, clinical research, consumer attitudes, and the broader technological environment that is shaping what tools will be available and sustainable in the near future.

Resource and Organization	Audience	Description and Direct Link

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Health Affairs: Artificial Intelligence in Health Care *Health Affairs*	Leadership · Policy	The leading health policy journal's curated AI topic page — peer-reviewed articles, commentary, and briefs on AI's impact on care delivery, payment, workforce, and access. Free and subscription content available. → healthaffairs.org/topic/ai
STAT News — Artificial Intelligence *STAT News*	All roles	Daily health and life science journalism with deep AI coverage — including regulatory developments, vendor news, real-world implementation stories, and access critiques. Recommended for leadership briefings. → statnews.com — AI Coverage
Digital Health and AI Funding Report (Annual) *Rock Health*	Leadership · IT leads	Annual intelligence on venture investment in digital health AI — useful for tracking which tools and vendors are gaining market momentum and which are losing traction before procurement decisions. → rockhealth.com — Annual Reports
AI Index Report (Annual) *Stanford Human-Centered AI Institute (HAI)*	Leadership · Strategy	The most comprehensive annual benchmarking of global AI capabilities, adoption, safety incidents, regulation, and economic impact. Provides a broader context for understanding where health care AI sits in the larger landscape. → aiindex.stanford.edu/report
JAMA Network — AI in Medicine *JAMA Network*	Clinical leadership	Peer-reviewed AI clinical research published across JAMA's flagship and specialty journals — covering diagnostic AI, predictive models, NLP in EHRs, and emerging guidance on AI validation methodology. → jamanetwork.com — AI in Medicine

AI Regulation and Policy

Federal and state regulatory requirements, guidance documents, and policy frameworks that govern AI in healthcare settings — including HIPAA implications, FDA oversight of clinical AI, compliance and emerging Massachusetts state-level legislation.

Resource and Organization	Audience	Description and Direct Link
HHS Artificial Intelligence Strategy and Policies *U.S. Department of Health and Human Services (HHS)*	Leadership · Compliance	HHS's official AI strategy, including guidance on responsible AI use across programs, algorithmic accountability requirements, and coordination with HIPAA compliance frameworks. Updated as federal policy evolves. → hhs.gov — AI Strategy
Executive Order on Safe, Secure, and Trustworthy AI — Health Provisions *White House / HHS*	Leadership · Compliance	Section 8 of the 2023 Executive Order on AI directed HHS to develop an AI safety program for health care, establish reporting mechanisms for AI-related harms, and advance access and affordability in AI deployment. → whitehouse.gov — AI Executive Order

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

FDA Artificial Intelligence and Machine Learning Action Plan *U.S. Food and Drug Administration (FDA)*	Clinical leadership · IT leads	FDA's framework for regulating artificial intelligence/machine learning-based (AI/ML-based) software as a medical device (SaMD). Critical reference when evaluating any AI-based clinical decision support tool subject to FDA oversight. → fda.gov — AI/ML Medical Devices
ONC Final Rule: Health Data, Technology, and Interoperability (HTI-1) *Office of the National Coordinator for Health IT (ONC)*	IT leads · Compliance	The 2024 HTI-1 rule introduces transparency requirements for decision support interventions (DSIs), including AI tools used in clinical workflows — directly affecting EHR-integrated AI tools in FQHC settings. → healthit.gov — HTI-1 Rule
State AI Legislation Tracker *National Conference of State Legislatures (NCSL)*	Leadership · Compliance	Real-time tracker of state-level AI legislation across all 50 states — including health care AI bills, data privacy requirements, algorithmic accountability laws, and consumer protection measures. → ncsl.org — AI Legislation Tracker
HIPAA and Artificial Intelligence: Guidance and FAQs *HHS Office for Civil Rights (OCR)*	Compliance · IT leads	HHS OCR guidance on how HIPAA Privacy and Security Rules apply to AI tools that use, generate, or share protected health information — including de-identification standards, BAA requirements, and breach considerations. → hhs.gov — HIPAA Guidance

Relevant Literature

The literature is robust discussing implementation, best practices, lessons learned and impacts of AI utilization in healthcare.

Resource and Organization	Audience	Description and Direct Link
Sharing a Workplace: How Health Systems Make Artificial Intelligence Useful to Clinicians *The Commonwealth Fund*	Clinical leadership · IT leads · Operations	Explore how healthcare staff are working to make AI useful in solving business problems and improving care. Identify strategies to mitigate the risk of harm to patients and healthcare institutions. → The Commonwealth Fund — Mitigating risk with AI implementation
The Missing Ingredient in Health Care AI? Community Voices *The Commonwealth Fund*	Clinical leadership · IT leads	*The Dose* podcast episode discussing what responsible AI innovation in healthcare looks like. Ensuring that by building AI tools for the most vulnerable patients, it will create a tool that works well for everyone. → The Dose — Patient voice in responsible AI development
Application of Artificial Intelligence in Community-Based Primary Health Care: Systematic Scoping Review and Critical Appraisal		*https://www.jmir.org/2021/9/e29839/*

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Building Healthier Communities: The Intersection of AI, IT, and Community Medicine		*https://asrjs.com/pdfs/building-healthier-communities-the-intersection-of-ai-it-and-community-medicine-160.pdf*
Bridging the Gap: From AI Success in Clinical Trials to Real-World Healthcare Implementation—A Narrative Review		*https://www.mdpi.com/2227-9032/13/7/701*

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Appendix E: AI Literacy Resources

Free and low-cost AI literacy training for community health center teams

This appendix compiles free and low-cost AI literacy training that community health center teams can use to build the knowledge the Playbook assumes. Resources span foundational AI literacy, healthcare-specific training, community health worker–focused training, and nonprofit AI adoption, organized by audience and cost tier. They support the AI literacy work that begins in Phase 1 — Listen and Learn (Section 3). Links are current as of June 2026 and should be re-checked before each publication.

Tier 1: Built for Community Health Workers and Healthcare Staff

Designed for frontline health workers, clinical staff, and healthcare administrators — the highest-priority options for community health center teams.

AI Literacy for Community Health Workers — TechChange and Johnson & Johnson

Co-developed with community health workers from organizations including the Community Health Impact Coalition, Living Goods, PATH, Dimagi, Medic, and the Gates Foundation. One of the most contextually relevant AI literacy tools for community health center staff, available globally with no paywall.

Format	Recorded video / self-paced online
What it covers	Defining AI; AI in everyday health work; data privacy for client information; advantages and limits of AI; responsible AI use alongside human judgment and community trust
Cost	Free (open access, no login)
Time	About 2 hours; ongoing (published February 2026)
Link	chwcentral.org/training-resources/artificial-intelligence-literacy-for-community-health-workers/

HelloAI — GE HealthCare (Foundations and Professional tracks)

More than 20 hours of healthcare AI education at no cost, built specifically for healthcare across clinical, engineering, research, administrative, and leadership roles. Includes expert-led sessions, journal clubs, and case studies.

Format	Self-paced on-demand learning plus recorded live webinars

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

What it covers	AI fundamentals, machine learning, generative AI, responsible AI; clinical and administrative applications; data privacy and security; deploying AI in practice
Cost	Free (both tracks)
Time	About 1 hour (Foundations); 20+ hours (Professional); ongoing (announced May 2026)
Link	helloaiprofessional.com

AI in Healthcare — RCSI University of Medicine and Health Sciences (Microsoft partnership)

Open to all healthcare professionals, managers, administrators, and allied health staff with no prior AI experience. Led by consultant radiologist Dr. John Sheehan.

Format	Self-paced online; 4 modules with quizzes and reflective exercises
What it covers	Core AI concepts and applications; ethics, safety, and governance; clinical applications; administrative applications
Cost	Free (optional CPD certificate with 10 CPD points, €50)
Time	About 10 hours; ongoing (enroll immediately)
Link	rcsi.com/online/find-a-course/cpd/a/i/ai-in-healthcare

Generative AI for Healthcare — Google Skills

Introductory course for clinicians, nurses, administrators, and researchers, with a shareable badge. Covers improving diagnostics, streamlining workflows, and tailoring patient care plans.

Format	Self-paced online; videos plus hands-on exercises
What it covers	Generative AI fundamentals; large language models in healthcare; clinical and operational use cases; prompt engineering
Cost	Free
Time	About 1 hour; ongoing (updated April 2026)
Link	skills.google/course_templates/1081

Discover AI for Leaders in Healthcare — Microsoft Learn

For healthcare business leaders: how to adopt AI within a health organization, including strategy, governance, and health-specific use cases. Part of the Microsoft Learn AI Hub.

Format	Self-paced online learning path; 5 modules
What it covers	AI adoption strategy; responsible AI practices; healthcare use cases; Microsoft AI resources for health systems

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Cost	Free
Time	Variable; ongoing
Link	learn.microsoft.com/en-us/training/paths/discover-microsoft-ai-leaders-healthcare-industry/

Tier 2: Nonprofit and Mission-Driven

Tailored to nonprofit professionals and social-impact teams, and directly applicable to community health center administrative, operational, and leadership staff.

Unlocking AI for Nonprofits — NetHope and Microsoft Elevate

Built for nonprofit professionals with practical tools and real-world use cases. Take one pathway or the full series, with an AI readiness assessment to guide planning. Over 5,000 enrollments as of June 2026. This content is currently available to existing members and partners.

Format	Self-paced; 4-pathway series (Kaya platform)
What it covers	AI basics; applications of generative AI (reporting, data, donor communications); Microsoft Copilot and advanced tools; responsible use of AI (ethics, responsible use, organizational safeguards)
Cost	Free (CPD-certified)
Time	A few hours per pathway; ongoing
Link	nethope.org/programs/unlocking-ai-for-nonprofits/

AI Skills for Nonprofits — Data.org and Microsoft

For nonprofit leaders and professionals with minimal AI experience, with customizable pathways. Completion earns a digital badge and can unlock micro-grants and team certification bonuses.

Format	Self-paced online; 3 modules
What it covers	Demystifying AI; crafting effective prompts; ethical decision-making for nonprofits; practical applications for mission-driven work
Cost	Free
Time	Variable; ongoing
Link	data.org/resources/ai-skills-for-nonprofits/

Tier 3: Universal and Foundational

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Open to any worker regardless of sector, requiring no prior AI experience, and useful for broad staff onboarding into AI literacy.

Make America AI-Ready — U.S. Department of Labor

Designed for maximum accessibility: any cell phone works, including basic flip phones, with no Wi-Fi, app, or laptop required. Aligned to the Department of Labor five-pillar AI Literacy Framework, and an ideal first step for front-desk and clinical support staff.

Format	Text message-based (10 minutes/day for 7 days)
What it covers	AI principles and core concepts; AI tools and use cases; prompt writing; evaluating AI outputs; responsible and ethical use
Cost	Free
Time	10 minutes/day for 7 days; ongoing (launched March 2026)
Link	Text READY to 20202, or dol.gov/ai-ready

IBM SkillsBuild — AI Fundamentals and Generative AI tracks

Over 1,000 courses in 20 languages, with AI fundamentals, generative AI, and AI ethics available free. Guided cohorts and self-paced access; partners with organizations serving underserved communities.

Format	Self-paced online (1,000+ courses; multiple AI pathways)
What it covers	AI fundamentals; generative AI; prompt writing; machine learning; AI ethics; AI customer service; chatbots; data science
Cost	Free
Time	Variable; ongoing
Link	skillsbuild.org

Microsoft AI Learning Hub (Microsoft Learn)

Curated learning paths for business and technical roles, with free skills badges. Includes Microsoft 365 Copilot training, which many community health centers may already have through nonprofit licensing.

Format	Self-paced online modules, learning paths, and badges
What it covers	AI basics; Microsoft 365 Copilot; Azure AI; responsible AI; AI agents; organizational AI transformation
Cost	Free
Time	Variable; ongoing
Link	learn.microsoft.com/en-us/ai/

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

AI for Everyone — DeepLearning.AI (Andrew Ng) via Coursera

One of the most widely recommended non-technical AI literacy courses, with over 7 million learners. Designed for non-STEM professionals; covers building AI strategy and collaborating with AI teams.

Format	Self-paced online; 4 modules (videos, readings, quizzes)
What it covers	What AI is; practical capabilities and limits; AI project planning; building AI strategy; AI ethics and societal impact
Cost	Free to audit ($49 certificate; financial aid available for a free certificate)
Time	About 6 hours; ongoing
Link	deeplearning.ai/courses/ai-for-everyone

Massachusetts AI Hub — Google AI courses (MassTech)

Through a Massachusetts partnership with Google, AI courses are available free to Massachusetts residents through December 31, 2027 — directly accessible to MLCHC member health center staff.

Format	Self-paced online (Google AI certificate tracks)
What it covers	Foundational AI skills and Google AI certificate pathways
Cost	Free for Massachusetts residents
Time	Available through December 31, 2027
Link	aihub.masstech.org/google-certificates

Tier 4: Emerging (Watch List for MLCHC)

AI in Healthcare for the Safety Net — CHAI and NACHC (in development)

The Coalition for Health AI and the National Association of Community Health Centers announced a partnership in August 2025 to co-design a curriculum tailored to community health centers and low-resource safety-net providers, plus an inaugural community health center survey to identify literacy gaps. One of the first purpose-built for the FQHC context once released.

Format	Virtual, hybrid, and in-person (multi-level curriculum in development)
What it covers	AI literacy and competency for community health center leaders, providers, and clinical staff at various levels; responsible AI use in safety-net settings; vendor evaluation for community-health-center-specific tools
Cost	To be determined (expected free or low-cost for CHC members)
Time	Not yet released
Link	nachc.org (follow for updates)

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Summary at a Glance

#	Resource	Audience	Format	Cost	Time
1	TechChange / J&J — AI Literacy for CHWs	Frontline CHWs	Video, self-paced	Free	~2 hrs
2	GE HealthCare HelloAI	All healthcare staff	On-demand + webinars	Free	1–20+ hrs
3	RCSI AI in Healthcare	Healthcare pros, administrators	Self-paced, 4 modules	Free*	~10 hrs
4	Google — Generative AI for Healthcare	Healthcare professionals	Self-paced + labs	Free	~1 hr
5	Microsoft Learn — AI for Healthcare Leaders	CHC leadership	Self-paced, 5 modules	Free	Variable
6	NetHope — Unlocking AI for Nonprofits	All nonprofit staff	Self-paced, 4 pathways	Free	A few hrs
7	Data.org — AI Skills for Nonprofits	Nonprofit leaders, staff	Self-paced, 3 modules	Free	Variable
8	DOL — Make America AI-Ready	All workers (any phone)	Text, 7-day	Free	10 min/day
9	IBM SkillsBuild	All staff	Self-paced + cohorts	Free	Variable
10	Microsoft AI Learning Hub	All staff	Self-paced	Free	Variable
11	DeepLearning.AI — AI for Everyone	Non-technical staff	Self-paced, 4 modules	Free**	~6 hrs
12	Massachusetts AI Hub — Google courses	All MA staff	Self-paced	Free (MA)	Thru 2027
13	CHAI + NACHC — Safety Net curriculum	CHC leaders, providers	Virtual/hybrid/in-person	TBD	TBD

**RCSI optional CPD certificate €50. **DeepLearning.AI free to audit; financial aid available for a free certificate. Full links appear in each entry above.*

Recommended Deployment Strategy

Frontline community health workers and clinical support staff: start with the TechChange/J&J course for its direct community-health-worker context, then the Department of Labor Make America AI-Ready text course as an accessible daily touchpoint.

Clinical and administrative staff: prioritize GE HealthCare HelloAI Foundations for its healthcare framing and the RCSI course for deeper clinical and governance grounding with CPD credits.

Operational, communications, and leadership staff: the NetHope and Data.org nonprofit series provide the most mission-aligned framing for non-clinical roles.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.

Executives and C-suite: Microsoft Learn AI for Healthcare Leaders and DeepLearning.AI's AI for Everyone provide strategic frameworks without requiring a technical background.

Watch closely: the CHAI and NACHC Safety Net curriculum is the only resource being co-designed specifically for FQHCs and safety-net community health centers, and it should be adopted as a priority training asset once released.

© 2026 Massachusetts League of Community Health Centers. Prepared by Phicil-itate Change LLC.