

Implementation Playbook

*Ambient AI: a practical,
vendor-agnostic,
evidence-based playbook
for health systems*






TABLE OF CONTENTS

03	What Does This Playbook Include?
03	Background & Purpose
04	Who Is This Playbook For?
04	How Is This Playbook Organized?
04	Listening In: A Summary of Challenges & Insights
08	Best Practice Statements
29	Appendix 1: Supporting Resources
30	Appendix 2: AI Lifecycle
32	Appendix 3: AI Governance Framework
34	Appendix 4: Consensus Process
35	Appendix 5: Thank you & Contributors
36	Appendix 6: Document Process & Governance



Ambient AI: Implementation Playbook

For Implementers

What Does This Playbook Include?

The Ambient AI Implementation Playbook is a practical, vendor-agnostic, evidence-based resource that aggregates consensus-defined best practices to enable responsible, system-level rollout of Ambient AI technologies. It is intended to support Implementers (i.e., health systems) in operationalizing AI lifecycle management and governance as they procure, deploy, and monitor ambient AI technologies.

Background & Purpose

Health systems of every size and type—from pediatric practices, community health centers, medium and large health systems, and academic medical centers—are rapidly adopting ambient AI technologies to ease documentation burden and restore clinician focus to the patient. Because this adoption is moving faster than shared standards can keep up, a multi-stakeholder work group was convened to bring these diverse organizational perspectives together in one place. Participants shared real-world implementation experience, surfaced practical insights, and identified the gaps and opportunities that cut across different care settings. Key ambient AI vendors were active participants throughout, working alongside health systems as the group drafted best-practice statements and defined them by consensus. The result reflects the combined perspective of a broad and varied community—spanning organization types and vendor viewpoints—working to translate early, fast-moving adoption into shared, durable guidance for the ambient space.

This Implementation Playbook is solely for the Implementer persona — health systems that are procuring, piloting, deploying, and monitoring ambient AI technologies produced by third-party vendors.

Primary Stakeholders

- Health system implementation teams; Clinical informatics and digital health leadership; Risk, safety and compliance officers; Clinicians using or piloting ambient AI solutions; Education leaders developing AI competency curricula; Vendor technical teams.

Note

While current Ambient AI adoption is most commonly associated with physician and advanced practice provider documentation workflows, future implementations (and iterations of this Implementation Playbook) may extend to nursing, care coordination, therapy services, case management, and other healthcare professions.

Who Is This Guide For?



Implementers: individual(s) at health systems responsible for the procurement, deployment, and/or overall realization of a system or component in accordance with a specified design (derived from IEEE 829 and IEEE 730)

How Is This Playbook Organized?

Sections *Listening In: A Summary of Challenges & Insights* and *Best Practice Statements* are organized by AI lifecycle step. The *Best Practice Statements* section is also tagged with applicable domain/subdomain(s) from CHAI's AI Governance Framework. Additional detail about the AI lifecycle and AI Governance Framework are included in appendices 2 and 3, respectively.

Each Best Practice Statement is listed under the earliest applicable AI lifecycle step. When a statement applies to multiple steps, its full AI lifecycle scope is noted in the AI lifecycle metadata so Implementers can carry the guidance through downstream implementation tasks.

As you navigate the *Best Practice Statements* section of the Implementation Playbook, a few practical suggestions are included here:

- For health systems just starting their ambient AI implementation process, focus on Procurement and Pre-Deployment best practice statements.
- For health systems piloting solutions, focus on Piloting best practice statements.
- For health systems scaling enterprise-wide ambient AI, focus on Deployment and Monitoring best practice statements.

Listening In: A Summary of Challenges & Insights

This section summarizes key challenges and insights related to ambient AI implementation, organized by AI lifecycle step.

Challenges/Insights – Procurement

Several procurement challenges center on the financial and access models that govern ambient AI. Per-user billing models can create unexpected financial exposure during deployment, particularly when user adoption accelerates faster than financial tracking or contract awareness, and standard “one-license-per-seat” pricing models are often incompatible with the reality of modern healthcare staffing, which includes part-time providers, residents, and contracted groups. Beyond cost, friction in the path to access—such as license-by-license provisioning, mandatory pre-use learning modules, and hard-to-find SharePoint pages—can stall adoption after the early-adopter wave; health systems that have moved fastest treat the licensing model, the provisioning workflow, and the training experience as a single design problem, where each touchpoint is designed to lower the activation energy required of a busy clinician. The feature landscape compounds these decisions, because the ambient AI feature set is shifting and feature limitations are one of the primary barriers to recruiting clinicians whose workflows depend on highly structured documentation. Health systems that treat the vendor relationship as an ongoing partnership are typically able to access new functionality faster and shape it to local needs, though certain feature categories—such as level-of-service or coding suggestions, order entry, and structured clinical actions—carry materially different compliance, billing, and patient-safety risk profiles than core transcription, and therefore warrant a different review pathway before they are released to the enterprise.

A second cluster of challenges concerns data governance and legal review. Legal and compliance reviews for ambient AI often hinge on where patient data is stored and processed, including whether vendor infrastructure resides outside the continental United States, while vendor contracts frequently contain ambiguous language around the use of de-identified patient data for model training and other commercial purposes. Consent adds further complexity, as patients may initially consent but decide to opt out during the session, particularly when the conversation shifts to sensitive topics like behavioral health; healthcare providers must be prepared to respond to mid-encounter withdrawal of consent, and health systems must have clearly defined workflows for handling both partial and full data deletion requests, including those that extend to vendor-held data. These concerns are heightened in behavioral and mental health encounters, which can generate particularly sensitive audio data—disclosures related to trauma, substance use, gender identity, and other protected information—that may warrant heightened data governance protections beyond what is standard for general clinical use cases. The evolving regulatory landscape, such as third-party data sharing agreements with government agencies, further amplifies the risk of uniform, one-size-fits-all data retention policies, and standard retention periods and a lack of configurable controls can cause entire clinical programs to opt out of ambient AI.

Finally, two cross-cutting risks shape procurement decisions. Unlike traditional software, ambient AI technologies are subject to performance drift—output changes that may stem from underlying model updates, shifts in usage patterns, evolving data, or policy changes—which can disrupt clinician workflow, documentation quality, and trust. And when enterprise-sanctioned solutions are unavailable, difficult to use, or not yet financially feasible, clinicians may turn to unvetted, off-the-shelf ambient technology, creating significant security and privacy risks for the organization and its patients.

Challenges/Insights – Pre-deployment

A substantial set of pre-deployment challenges are legal and consent-related. The fundamental legal challenge with ambient AI is not the artificial intelligence itself, but the act of capturing a patient's literal voice, which triggers diverse and often stringent state-level wiretapping and eavesdropping statutes; because the legal definition of "recording" is still evolving in the context of transient audio processing—where audio is used to generate a note and then immediately deleted—a conservative approach is important to mitigate liability across different jurisdictions. Standard HIPAA notices and general intake disclosures may be insufficient to satisfy consent requirements for ambient AI because these tools passively record healthcare clinician-patient conversations and transmit audio or transcripts to third-party vendors for processing, creating consent obligations that extend beyond traditional documentation software (including potential exposure to state wiretapping statutes and confidentiality protections that vary significantly by jurisdiction). There is also a distinction between a legal obligation and an ethical one, and while legal requirements for recording vary, relying on a general Notice of Privacy Practices may meet legal standards but can miss opportunities to build patient trust. Consent is especially complex in pediatric settings, where encounters routinely involve multiple speakers (parents, guardians, caregivers, siblings), adolescents with confidential health concerns (such as reproductive health, mental health, and substance use), and jurisdiction-specific recording consent requirements; a one-size-fits-all consent process will fail to address these realities and exposes organizations to legal, ethical, and reputational risk, so organizations serving pediatric populations should design consent workflows that account for these unique dynamics from the outset. Documentation status raises a related risk, because AI-generated documentation is a draft requiring active clinician review and is not automatically accurate or complete, and the concurrent existence of an AI-generated transcript and a finalized medical record creates a "dual record" dynamic that can expose health systems to legal risks if audio and transcript data are not governed under a clearly defined retention and record status policy.

How the technology is framed and tailored also matters before deployment. Establishing a productivity-linked, individual-contributor Return on Investment (ROI) model can create a "productivity tax" perception that undermines one of the technology's key objectives of reducing clinician burnout; when ambient AI is framed as an optional efficiency tool that providers must subsidize through increased patient volume, it can hinder adoption, whereas treating the technology as a core enterprise utility—comparable to an email system or the EHR—aligns the investment with broader organizational goals for workforce stability, recruitment, and professional satisfaction. Tailoring is equally important, because ambient performance and clinician adoption vary significantly by specialty, and a generic approach to documentation templates and training often fails in high-acuity environments like the Emergency Department.

Challenges/Insights – Piloting

Several piloting insights focus on how performance is measured and made visible. Reviewing only the final AI-generated clinical note obscures the frequency, magnitude, and downstream impact of safety risks such as hallucinations and omissions, limiting a health system's ability to meaningfully monitor performance over time. Aggregate or generic vendor performance data compounds this blind spot, since it may mask meaningful variation in accuracy, omissions, and usability across clinical specialties and care team roles, particularly in high-complexity domains or those with specialized terminology (such as imaging) where expectations and standards are not yet well defined. The stakes of getting

measurement right are rising, because as ambient AI capabilities expand toward suggesting orders, flagging diagnoses, and supporting level-of-service determinations, the clinical and safety stakes of AI-generated documentation increase significantly—yet most organizations currently measure ambient AI success through adoption rates and clinician satisfaction, without systematically evaluating how AI-generated notes influence downstream impacts.

Piloting also surfaces challenges specific to particular environments. Using ambient AI in a teaching environment introduces unique challenges for the implementing organization regarding educational benchmarks, supervisory accountability, and technical ownership of clinical documentation, and these challenges vary by learner type, discipline (such as medicine, nursing, and advanced practice), and the technical configuration of the EHR. Pediatric environments present unique audio challenges, including high ambient noise (such as crying children) and the presence of multiple caregivers, which requires intentional verbal cues from the clinician to ensure accurate documentation. Consent requirements and workflows likewise vary significantly between clinical settings (for example, inpatient nursing versus outpatient provider visits), across specialties, and by geographic location, which requires tailored but standardized documentation methods; health systems must also determine whether to implement a holistic consent strategy across all specialties, unique department-level consent workflows, or a hybrid of both.

Finally, who gets access during piloting shapes where value emerges. Centralized administrative planning often fails to anticipate where the most significant clinical efficiencies or novel applications of ambient AI will occur, and when access is restricted based on predicted utility, health systems risk stifling “bottom-up” innovation driven by clinicians who understand the unique documentation burdens of their specific workflows. Ultimately, successful enterprise adoption relies on aligning organizational requirements for measurable efficiency with the clinical desire for a more sustainable and effective practice environment.

Challenges/Insights – Deployment

Large-scale rollouts, whether net-new or replacement, can create financial and operational strain if not aligned to the organization’s transition context, size, and resource base, and the choice between staged and “big bang” rollouts involves trade-offs that should be weighed through the organization’s governance structure. Ease of activation can be deceptive, because ambient AI can be easy to activate but not automatically safe or effective; intuitive interfaces reduce onboarding friction, yet appropriate use still depends on structured guidance around workflow, consent, and human oversight. Success ultimately depends on bidirectional adaptation—clinicians adjusting their communication style and understanding the technology’s technical requirements within the documentation system (the EHR, or specialty subsystems such as CVIS or RIS), and the technology aligning with the prevailing workflows, documentation formats, and communication styles of the clinician user groups it serves.

Consent and recording remain central at deployment, and requirements are subject to highly variable state laws and institutional risk tolerances. Standardizing encounter-level verbal consent, supported by written opt-out safeguards, builds trust in ambient AI use and must align with state- and age-specific consent requirements. Passive notification, such as signs in a waiting room or mentions in upfront paperwork, may be missed by patients and fails to provide affirmative confirmation of AI use; however, passive notification may be legally sufficient in one-party consent states, and many health systems

prefer this approach as less intrusive than active consent processes. The more critical factor is understanding your region’s regulatory requirements, applicable consent laws, and patient population.

Monitoring and education round out the deployment challenges. Relying solely on subjective provider feedback may overlook emerging risks or inequitable impact; however, meaningful monitoring requires metrics that are context-specific and defined in advance. On training, highly prescriptive curricula can create significant impediments to rapid enterprise-wide adoption, and because many clinicians demonstrate a high degree of autonomy in learning the fundamentals of ambient technology, effective educational strategies may prioritize agile, “just-in-time” support structures over rigid coursework while firmly establishing the “human-in-the-loop” paradigm, ensuring clinicians view the technology as a supportive assistant rather than a replacement for professional clinical judgment. Adoption is also not uniform—users span a spectrum from early innovators to skeptical laggards, and a single training modality may inadequately serve the full population—so organizations may consider offering flexible support pathways matched to user needs, calibrating training requirements to balance readiness with barriers to entry as familiarity matures over time, and maintaining active community channels for peer knowledge sharing. Even so, once early adopters have been captured, adoption sometimes plateaus, and two patterns consistently break that plateau in health systems: communication that meets clinicians inside the systems they already use, segmented to the specific persona or specialty receiving the message; and monitoring dashboards whose metrics are real-time and actionable, segmented by setting of care, persona, and use pattern so that operations teams can identify low-utilization pockets and intervene with the right message to the right audience.

Challenges/Insights – Monitoring

Sustaining ambient AI adoption requires continuous monitoring of both technical performance and the technology’s impact on clinician well-being and patient satisfaction. This is complicated by the fact that ambient AI technologies are dynamic and vendor-controlled, meaning software updates or backend changes can occur without notice and potentially alter the technology’s performance over time. Because many health systems lack the infrastructure for sophisticated in-house monitoring, proactive vendor communication strategies are an equally important safeguard alongside internal tracking mechanisms.

Best Practice Statements

Each Best Practice Statement is organized by AI lifecycle step. Each statement is also tagged with its applicable domain/subdomain(s) from CHAI’s AI Governance Framework. Additional detail about the AI lifecycle and AI Governance Framework are included in appendices 2 and 3, respectively.

As a reminder, each Best Practice Statement is listed under the earliest applicable AI lifecycle step. When a statement applies to multiple steps, its full AI lifecycle scope is noted in the AI lifecycle metadata so Implementers can carry the guidance through downstream implementation tasks.

AI Lifecycle Key



 Stage(s) this best practice applies to

 Not applicable

AI Governance Framework Key

CHAI AI Governance Domains



AI Policy



Organizational
Structures



Organizational
Resources



Responsible AI
Lifecycle
Management & Use



Risk & Impact
Assessments



Responsible Data
Management & Use



Third Party
Management



Education,
Training &
Feedback

Procurement


BP1: Vendor billing structure and internal cost tracking



Organizational Resources

Clarify billing structure, activation timing, and month-end reconciliation processes with the vendor, and establish internal tracking mechanisms to monitor active users against contracted thresholds in order to prevent unanticipated cost escalation.

BP2: Vendor partnership for post-deployment monitoring and model-change review



Third Party Management

Screen vendors for willingness to partner on post-deployment monitoring, including providing roadmaps, advance notification of material model or system changes, and clinician-understandable release notes. During Monitoring, implementers and vendors should treat each other as partners, with recurring operational meetings (e.g., quarterly cadence) covering utilization, issue tracking, performance trends, and roadmap items. Major ambient model and/or system updates (e.g., underlying LLM changes,


significant retraining, or alterations to summarization, tone, or documentation structure) should trigger a formal review process similar to initial procurement vetting, including security re-review, where applicable.

Examples

One health system runs a 1-hour standing meeting with each vendor every 2 weeks, attended by Leadership, IS, and EHR analysts. The agenda covers total licenses in use, utilization by clinician, license distribution by care setting (Ambulatory, ED, Inpatient), an issue tracker, new feature releases, and forward roadmap items.

As another example, a large academic health system established contractually required quarterly operational reviews attended by clinical informatics leadership, IT security, compliance, and frontline clinician representatives. The contract required advance notification of any material model updates and clinician-understandable release notes, positioning the vendor as a trusted partner.

BP3: Vendor data privacy/deletion processes and mid-encounter withdrawal-of-consent workflows



Verify that vendors maintain clearly documented privacy and data governance processes, including mechanisms for off-cycle data access, correction, and deletion requests. Health systems should ensure these processes are easily accessible to administrators and that vendors can fulfill such requests within legally required timelines as defined by local regulatory, legal, and institutional governance requirements. Note that off-cycle deletion capabilities may carry additional cost from vendors; health systems should proactively evaluate and negotiate this during the procurement/contracting phase. Healthcare providers should also be trained on how to handle intra-encounter withdrawal of consent, with workflows defined for two distinct scenarios:

1. Partial withdrawal (patient does not want a specific statement captured) – the healthcare provider should pause the recording for that portion of the conversation and resume only when the patient is comfortable, or manually exclude that content from the generated draft note before it is finalized.
2. Full withdrawal (patient wants the entire encounter deleted) – the healthcare provider should stop the recording, delete the draft note in the EHR, and immediately initiate the defined process for notifying the vendor to delete any corresponding audio or transcript data stored on their systems. Health systems should establish this workflow with vendors in advance, either through direct deletion within the ambient AI application itself or via a defined rapid vendor escalation process, so that it is executable at the point of care.

Example

For example, a patient consents to the use of ambient AI documentation at the start of a clinic visit. Ten minutes into the encounter, while discussing a sensitive topic (e.g., mental health or substance use), the patient states they no longer feel comfortable being recorded and asks that the recording stop and that anything already captured be deleted. The healthcare provider stops the recording immediately and deletes the draft note generated in the EHR. However, because the ambient AI vendor may have stored portions of the encounter audio or transcript on their system for processing, the health system must have a mechanism to notify the vendor and request deletion of the corresponding audio or transcript data associated with that encounter.

BP4: Data residency review



Conduct a rigorous legal and information security review that specifically addresses data residency requirements. Confirm whether vendor infrastructure stores or processes patient data outside the continental United States, and assess alignment with applicable state regulations and organizational privacy standards.

Example

One organization experienced significant delays in ambient AI deployment because state Medicaid regulations prohibited billing for clinical notes if the underlying data had been accessed or stored outside the continental United States. Leadership emphasized the importance of confirming data residency specifics with vendors before contract execution to avoid downstream compliance failures.

BP5: Contractual limits on patient data use for model training




Request vendors to explicitly define how patient data may be used beyond direct care delivery. Specifically, contract language should clarify whether patient data, including de-identified data, will be used to train proprietary AI models or for other commercial purposes. Organizations should ensure these terms align with organizational privacy standards and applicable regulations before contract execution.

Example

One organization's leadership noted that vendor contracts often include broad rights to use de-identified data for model improvement. Without precise contractual definitions, organizations risk unknowingly contributing patient data, even in de-identified form, to proprietary model training.

BP6: Flexible licensing models for a pluralistic workforce



Consider flexible licensing models such as utilization-based pricing or specific part-time/resident tiers to ensure the technology can be cost-effectively deployed across a pluralistic workforce.

Example

One organization established utilization-based pricing to manage locums, part-time staff, and providers on leave (FMLA), ensuring the organization only pays for active use across various sites.

Note that account onboarding can be significantly more complex for clinicians who do not use the hospital's main email domain.

BP7: Reduce "shadow AI" through policy, enterprise-vetted alternatives, and traffic monitoring



Reduce unauthorized ambient AI use by combining (1) clear, well-communicated policy on permitted and prohibited tools; (2) workflow-integrated, enterprise-vetted solutions where feasible (e.g., embedded in the mobile EHR application); and (3) active monitoring of network traffic for unauthorized ambient usage (as able). Where a fully vetted enterprise solution is not yet financially or operationally feasible, organizations should still publish explicit guidance on disallowed consumer tools.

Example

IT security teams at one organization monitored traffic to consumer ambient AI websites and used the resulting data to identify gaps in their sanctioned tooling and accelerate procurement of an enterprise-vetted alternative. Integrating that solution directly into the EHR's mobile application was a key strategy to keep clinicians within the organization's secure ecosystem rather than resorting to personal devices or unvetted web-based solutions.

BP8: Context-specific data governance for behavioral and mental health encounters



Establish context-specific data governance policies for behavioral and mental health ambient AI use cases, including:

1. Configurable audio and transcript retention periods. If able, offer the ability to apply shorter retention windows or opt-out-of-retention for encounters involving sensitive disclosures;
2. Third party vendor contractual requirements that prohibit use of de-identified data for model training without explicit organizational consent; and
3. Clearly documented clinician protocols for pausing recording during sensitive segments of a visit.

These requirements should be negotiated at the procurement stage and re-evaluated throughout the vendor relationship.

Example

One health system noted that ambient AI implementation in behavioral health settings surfaces unique data governance challenges not encountered in general clinical use. Standard 30-day recording retention policies were identified as a particular concern. The health system noted that some entire clinical programs serving sensitive populations had declined to adopt ambient AI entirely, citing concerns about data retention and the risk of third-party access to recordings. A clinician leader called for vendors to offer configurable retention controls (including a "do not retain" option) to make the technology accessible and appropriate for behavioral health contexts. The group also acknowledged the heightened sensitivity of ambient AI in settings where patients disclose information related to gender identity or substance use disorders, and the importance of proactively addressing these risks with vendors at the procurement stage.

BP9: Designing access, training, and provisioning to minimize clinician friction



Design access, training, and provisioning pathways for ambient AI to minimize clinician friction. This may include actions such as evaluating broader licensing strategies (such as enterprise-level access where feasible), automating provisioning by user role wherever possible, and offering training in multiple modalities (self-guided videos, quick-start guides, and at-the-elbow support) so clinicians can engage on their own terms.

Example

One multi-hospital academic system in the Northeast attributed a meaningful inflection in adoption to three concurrent changes: (1) shifting from ad-hoc, request-based provisioning to role-based access via a native EHR allow-list, which let any inpatient, emergency department, or ambulatory clinician begin using the tool with minimal administrative overhead; (2) lightening a previously mandatory training module that had become a barrier to entry, and replacing it with multiple entry points including peer-led, supervisor-led, and self-service training; and (3) operating under an enterprise-wide license that removed seat-by-seat economics from the rollout decision. The system reported that the two highest weeks of new user uptake since program launch occurred immediately after these changes were implemented in combination. A separate large academic medical center in the Southeast reached comparable scale through a different but equally deliberate path: offering four to six onboarding sessions per week, two to three weekly Q&A lunch hours, and on-demand 1:1 sessions, all surfaced through a single QR-code landing page promoted on screen savers and email blasts.

BP10: Formalized vendor partnership and tiered release pathway for higher-risk features



Formalize your vendor relationship as an ongoing partnership with embedded vendor representation, establish a tiered release pathway (e.g., alpha to clinician champions, beta to broader enterprise) with active feedback obligations from designated champions. Route higher-risk features through a defined compliance and risk-review pathway before broader release.

Example

One multi-hospital system entered a new phase of vendor partnership in which it now serves as a testing site for alpha-stage features. The organization elevated the requirements for its clinician champions: continued access to early features is contingent on active participation and feedback, with non-participating champions quietly transitioned out. New features flow through a short alpha evaluation by champions (using a shared channel for notifications and feedback), then transition to a beta release across the enterprise if no issues surface. Critically, features with potential compliance or billing implications (e.g., level-of-service suggestions) are routed directly to the compliance team, which defines a separate, longer review pathway baked into the partnership agreement. A separate large academic medical center reinforces this approach by piloting higher-risk new functionality (e.g., order entry) with a small set of clinicians ahead of a tentative go-live, using their feedback to shape the communications and training that will accompany broader deployment. Across both organizations, an embedded vendor success associate participating in regular implementation team meetings is cited as a key enabler of this model.

Pre-deployment

BP11: Incorporate baseline privacy, data ownership, and EHR integration requirements



Incorporate baseline requirements for information privacy, data ownership, and EHR integration into a public Request for Proposal (RFP) to ensure vendor alignment with institutional rules, as well as state and national legal requirements. RFPs also surface information such as pricing, product features, future roadmaps, and training requirements for onboarding (e.g., understanding what the typical training process looks like from a vendor). Organizations may act as "conservative institutions" regarding data, requiring vendors to agree to limited trial periods to test efficacy and safety before committing to a purchase. Integration with existing systems, such as the Electronic Health Record (EHR), is a critical requirement for long-term user adoption.

BP12: Cross-functional governance bridging clinical operations and clinical informatics



Establish a cross-functional governance structure that identifies joint responsibility between the person/team responsible for change management in clinical operations and the clinical informatics person/team. A joint structure ensures that responsibility for deployment, training, and handling of technical builds and performance reporting is established. While cross-functional collaboration is critical, a clear owner(s) or champion(s) is also needed. Additionally, Implementers should define and include their implementation timeline and style (e.g., slow rollout by department, big bang, etc.) and ensure all necessary colleagues are aligned.

Note

This can be at the organization or network-level and may consist of one to many employees, depending on resourcing and size of organization.

BP13: Explicit goals for patients, clinicians, and the health system



Explicitly define and document the intended goals of ambient AI adoption across three core stakeholder groups: patients, practitioners, and the health system. These goals should articulate anticipated benefits

(e.g., improved documentation quality, reduced provider burden, enhanced patient communication, operational efficiency), as well as potential tradeoffs and risks. Clear goal-setting enables alignment across governance, monitoring, contracting, and training decisions and ensures that success metrics are tied to meaningful clinical, ethical, and operational outcomes rather than adoption alone.

Example

One health system found it critical to focus on "augmented intelligence" where the provider stays in the loop, prioritizing the reduction of "pajama time" and cognitive load over simply increasing patient volume. By identifying these specific pain points early—such as high provider burnout and delayed note closure—the organization ensured the technology was viewed as a solution to a problem rather than an added burden.

BP14: Standardize hardware, device ownership, and minimum security standards



Determine whether additional hardware (e.g., microphones, smart devices, mobile phones) is required and clearly define device ownership (health system-issued vs. practitioner-owned). Minimum security, privacy, and technical standards should be established in advance, including mobile device management expectations where personal devices are used.

Example

One health system mentioned that ambient technology should ideally be bidirectional and work across both Android and iOS platforms to accommodate various provider preferences. In instances where certain mobile platforms have limited functionality, organizations have successfully provided "loaner" iPads or dedicated phones to ensure that all providers, especially those in high-volume areas like the Emergency Department, can utilize the tool effectively. This proactive approach also prevents the risk of staff using non-vetted, non-HIPAA-compliant cloud solutions.

BP15: Cross-functional governance including legal, risk, compliance, IT security, and ethics



Establish a cross-functional governance structure (if one is not already in place) that includes legal, risk, compliance, IT security, and ethics expertise. This can be with one or many individuals depending on the resourcing and size of the institution and may come at the individual health system or network level. The governance structure is required to vet solutions and define ethical deployment standards prior to system-wide rollout.

Example

One health system found it necessary to form a dedicated AI Governance Committee to prevent providers from independently adopting non-vetted, potentially non-HIPAA-compliant "shadow AI" solutions. Additionally, an AI Ethics Work Group was established to ensure that AI use remains inclusive, transparent, and respectful of patient rights across the entire health system.

BP16: Engage legal/privacy counsel to assess sufficiency of existing consent frameworks




Engage legal and privacy counsel to assess whether existing consent frameworks – including standard HIPAA notices and general intake disclosures – are sufficient for ambient AI technology use, or whether explicit, visit-level consent mechanisms are required. Legal and/or privacy counsel should specifically evaluate state wiretapping statutes, behavioral health confidentiality laws, and other jurisdiction-specific regulations, particularly where audio or transcripts are transmitted to and stored by third-party vendors.

BP17: Develop a policy on clinician review obligations and "dual record" governance



Develop and implement a comprehensive policy that governs healthcare provider obligations for AI-generated documentation review, clearly establishing that audio recordings and transcripts are transitory draft communications (i.e., not the legal medical record) until reviewed, edited, and approved by the healthcare provider. The policy should explicitly define the record status of AI-generated outputs at each stage of the documentation workflow.

BP18: Complement legal consent with transparent verbal notification



Complement formal legal consent (integrated into general privacy policies) with a brief, transparent verbal notification to the patient at the start of the encounter. For example, one organization's legal team determined that specific verbal consent was no longer required after the general privacy policy was updated to include AI. However, many pediatricians continued to use a "hand-washing" approach – mentioning they were using their phone to help them chart while performing clinical tasks – which built camaraderie and resulted in almost zero patient opt-outs in this scenario.

BP19: Conservative wiretapping/eavesdropping assumption regardless of audio storage



Until courts or legislatures provide clearer guidance, health systems should conservatively assume that the use of ambient AI documentation tools constitutes audio recording governed by state wiretapping or eavesdropping statutes. As such, organizations should implement consent practices and operational safeguards consistent with applicable recording laws in their jurisdiction, regardless of whether the vendor ultimately stores audio, converts it to transcripts, or processes the audio transiently.

Note

For all group settings, these safeguards should account for the parent or guardian as the primary consenting party, especially since some states may have specific regulations regarding the recording of minors.

BP20: Centralized funding aligned to wellness, recruitment, and retention



Establish a centralized funding and value-alignment strategy that prioritizes provider wellness, recruitment, and retention as primary success metrics. This explicitly decouples the procurement of the technology from individual-level productivity requirements or provider-funded fee structures.

Example

During the piloting phase at one large healthcare organization, leadership initially considered a model where the cost of the technology was associated to individual providers, with the expectation that they would cover the expense by seeing additional patients. Leadership ultimately rejected this approach, arguing that Ambient AI should be treated as an enterprise-wide infrastructure utility. This decision was validated by pilot data showing that 92% of participants reported reduced burnout and 93% reported improved work-life balance.

Furthermore, the organization found that the technology became a critical recruitment and retention tool, with prospective high-level clinical staff viewing ambient documentation as a "minimum" requirement for their professional environment.

BP21: Specialty-tailored documentation templates and training modalities





Tailor documentation templates and training modalities to the specific clinical workflows of each specialty, ensuring that support channels (e.g., chat platforms) match the communication habits of that specific user group.

Example

One organization found that while ambulatory providers utilized internal team chat platforms for real-time support, ED providers rarely checked those channels. The implementation team had to pivot to using the EHR's internal, secure chat to reach ED clinicians where they were already working. Additionally, generic documentation templates were found to be insufficient for specialty use, requiring analysts to create specialty-specific versions (e.g., for vascular surgery or oncology) to ensure the AI captured the necessary nuance.

BP22: Pediatric-specific consent workflows (adolescent rights, multi-speaker, recording laws)



Health systems serving pediatric populations should develop and operationalize ambient AI consent workflows that explicitly address:

1. Adolescent-specific consent rights and confidentiality protections;
2. Multi-speaker encounter dynamics;
3. Applicable state or regional recording consent laws; and
4. Clinician discretion protocols for pausing or stopping recording during sensitive conversations.

Consent processes should be documented in the clinical record and reinforced through mandatory clinician training that includes consent attestation prior to system access.

Example

One pediatric health system operating in a two-party recording consent state developed a comprehensive consent framework that included: general treatment consent language indicating AI might be used; verbal consent required from all parties present in the room; clinician-facing consent scripts and patient-facing educational handouts; and a documented consent statement appended to the clinical progress note. E-learning training included an explicit attestation confirming clinicians understood all consent requirements.

Piloting

BP23: Feasible monitoring that combines vendor reporting with internal workflow signals



Responsible AI Lifecycle Management & Use

Define and operationalize a feasible monitoring approach that combines vendor-provided reporting with internal workflow signals. Specifically, Implementers (health systems) might:

1. Request vendor reporting on note quality and accuracy across prioritized specialties and care settings, including known error patterns and safety mechanisms;
2. Clarify responsibility upfront for which metrics are supplied by the vendor versus monitored internally, based on available EHR and analytics capabilities;
3. Implement internal safety monitoring using signals the organization can reasonably access, such as provider note-editing behaviors, audit samples for accuracy, chart closure time, and EHR usage patterns;
4. Establish a monitoring cadence and escalation plan during implementation, aligning on how findings will be reviewed, shared, and acted upon over time.

BP24: Specialty-disaggregated clinician feedback and vendor reporting



Responsible AI Lifecycle Management & Use

Risk & Impact Assessments

Engage clinicians and care team members from specialties representative of the institution's clinical footprint where ambient technology is reasonably expected to be adopted, to gather structured qualitative feedback on documentation accuracy, omissions, and clinical workflow fit. Request specialty-disaggregated performance reporting from vendors (e.g., accuracy and omission rates broken out by specialty and care setting) rather than attempting gold-standard or objective error quantification internally. Vendors should transparently communicate known specialty limitations.

BP25: Learner access policies in teaching environments



Organizational Structures

Education, Training & Feedback

Define clear policies regarding learner access to ambient technology across applicable clinical training programs. This includes determining whether trainees receive licenses, under what supervisory conditions use is permitted, and how documentation attribution, accountability, and liability (with particular attention to medical and other clinical students) will be managed. Because learners often rotate across multiple clinical sites, guidance should clarify where the tool may be used, how site-specific consent and recording requirements apply, and how policy differences are communicated. Policies should account for workflow and attestation differences across learner types (e.g., medical students vs. residents/fellows, and analogous distinctions in nursing or advanced practice training) and align with

the technical configuration of the EHR. Where granular access controls by training level are not technically feasible without significant maintenance or security redesign, organizations may adopt a clinically governed model in which the supervising clinician determines learner use on a session-by-session basis.

Note

Some health systems have limited ambient AI use to senior residents (e.g., 3rd and 4th year) to ensure that junior learners first master traditional documentation skills.

It is important to differentiate residents/fellows v. medical students in this context because the way supervisors attest documentation for these two groups may differ. For example, if a clinician attested a resident/fellow note, the resident/fellow would own the documentation and the clinician would own the attestation. However, for a note by a medical student, the clinician would typically edit and attest the documentation (making the clinician the owner).

BP26: Integrate patient consent capture into the standard clinical workflow



Integrate patient consent capture into the standard clinical workflow – such as during the nursing admission assessment or at the start of an outpatient visit – utilizing discrete data fields or standardized phrases within the EHR. When designing the consent strategy, health systems must explicitly account for both specialty-specific and state/jurisdiction-specific requirements, as applicable laws may mandate written consent documentation rather than verbal consent alone.

Example

When determining the consent approach, health systems should assess whether a holistic consent process (e.g., embedded in standard intake paperwork) is sufficient across departments, or whether certain specialties require supplemental workflows. For example, behavioral health departments may require additional verbal or written consent steps depending on state or regional regulations, regardless of what is captured in general intake documentation.

For inpatient settings, documenting consent in a discrete data field allows for organizational transparency, whereas outpatient providers may find dot phrases (e.g., Epic SmartPhrase) within encounter notes more efficient for recording verbal consent. However, verbal consent captured via dot phrase may not satisfy all legal requirements – for instance, Illinois Public Act 104-0054 requires mandatory written consent documentation for AI scribe use in behavioral health contexts.

Prior to finalizing any consent workflow, health systems should conduct a thorough review of applicable state laws, as requirements vary significantly by jurisdiction and continue to evolve.

BP27: Active narration and verbal identifiers in pediatric environments



 Responsible AI Lifecycle Management & Use

Train pediatric clinicians to use active narration and verbal identifiers during the encounter to help the AI distinguish between multiple speakers, such as parents or guardians. In a pediatric setting, a provider might say, "Thanks, Mom, for that history," which allows the AI to correctly attribute the preceding statements to the mother in the generated note.

Note

For pediatrics, ambient implementation should include pediatric-specific templates and configurations. As an example, pediatrics use unique terminology and approaches that other specialties do not use such as mood trackers to determine children's pain (which is quite unique from adult).

Experience shows that current models are fairly effective at filtering out child-related background noise (e.g., screaming) while still capturing the relevant medical conversation

BP28: Democratized, enterprise-wide utility access for all documenting roles



 Organizational Structures Organizational Resources

Consider a democratized enterprise-wide utility model, providing Ambient AI access to all clinical and ancillary roles with documentation responsibilities rather than limiting licenses to specific high-volume specialties. For example, one large health system discovered that while leadership initially prioritized standard ambulatory care, the most enthusiastic adoption and creative use cases emerged from high-pressure environments like emergency medicine during peak volumes and specialized areas like pathology for autopsies. By providing broad access, the organization empowered "contextual masters" (e.g., frontline staff with the deepest understanding of their own workflows) to determine how to best integrate the technology into diverse settings.

BP29: Multifaceted evaluation framework (operational metrics + wellness + real-time feedback)



 Organizational Structures Organizational Resources

Establish a multifaceted evaluation framework for ambient AI that measures success through both operational metrics (e.g., same-day chart closure, coding accuracy) and provider wellness indicators, while utilizing real-time provider feedback to demonstrate total value to executive leadership.

Example

One organization achieved a 96% same-day chart closure rate and improved billing accuracy by capturing the full nuance of clinical conversations without the need to increase patient volume. To secure additional licenses, leadership presented unedited screenshots of enthusiastic provider feedback from internal message boards directly to the Board of Directors, effectively humanizing the operational data to demonstrate the technology's impact on the staff.

BP30: Systematic framework for monitoring downstream impact of AI-generated documentation



Establish a systematic framework for monitoring the downstream impact of ambient AI-generated documentation (e.g., spanning clinical decision-making, documentation completeness and accuracy, prior authorization denial rates, diagnostic variation, and patient-reported experience) and establish a formal risk and impact assessment reviewed by AI governance leadership, as organizations begin to adopt AI capabilities that directly influence clinical decisions, order generation, or diagnostic or billing coding. Monitoring metrics should be defined prior to deployment so that baseline comparisons are available to detect unintended consequences as AI scope expands.

Deployment

BP31: Evaluate staged vs. "big bang" rollouts via governance structure




Use organizational governance structure to evaluate the risk-benefit of staged versus big-bang rollouts based on size, financial capacity, and transition context. When a staged approach is selected, executing in waves (e.g., predefined blocks of licenses) allows early waves to serve as structured evaluation phases to test performance, safety, and workflow impact before broader deployment, while enabling financial offsetting and operational adjustments.

Example

For example, health systems with legacy human scribe programs can phase in ambient technology as those programs experience natural attrition, and starting with enthusiastic early adopters can build initial momentum. Under-resourced organizations without prior scribe programs may face a net-new cost burden and should explore funding mechanisms such as HRSA grants or group purchasing arrangements. In replacement ("rip-and-replace") scenarios where staged licensing is not feasible, focus on coordinated transition planning and structured monitoring to reduce disruption during the system-wide shift.

BP32: Flexible, role-appropriate onboarding addressing automation bias and consent



Ensure onboarding is flexible and role-appropriate – incorporating in-app guidance, brief self-directed modules, or hands-on support – while clearly addressing automation bias, patient consent requirements, data privacy expectations (including risks of shadow use), and product-specific workflow integration. Training should differentiate between providers new to ambient technology and those transitioning from another vendor, and should reinforce ongoing reinforcement rather than a one-time session.

BP33: Define a priori performance and impact indicators



Define a priori which performance and impact indicators are most relevant to your clinical context and operational goals, recognizing that appropriate metrics will vary across specialties and system size. These may span domains such as utilization patterns, note length, chart closure delays, documentation quality, equity considerations, and financial impact. Health systems should collaborate with vendors to establish feasible dashboards, clarify data sources and baselines, and set an agreed-upon cadence for reviewing trends over time.

Note

If the health system isn't capable of this level of sophistication, consider assessing whether the vendor themselves can provide this reporting.

BP34: Auditable patient consent capture and documentation process



Establish a clear and auditable process for obtaining and documenting patient consent (verbal or written) including defined workflows for how consent will be captured and displayed in the final medical record. Special protocols should address pediatric and minor consent requirements, which may differ by jurisdiction. Training should reinforce proper consent documentation practices, particularly in scenarios where consent is obtained verbally prior to recording, to ensure transparency and regulatory compliance.

BP35: Standardized per-clinician implementation checklist and configurable onboarding



Utilize a standardized implementation checklist for each clinician (or group of clinicians) that covers topics such as technical setup, documentation template modification (e.g., inserting "markers"), and training on verbal prompting to support high-quality output. Vendors should provide onboarding resources and support user-configurable documentation formats that reflect variation across organizations and specialties. Where the solution is sufficiently intuitive and the foundational technical work has been completed (e.g., validated vendor configuration, specialty mapping, properly inserted markers), organizations may offer flexible onboarding pathways (e.g., self-start materials, optional live training) rather than requiring a mandatory live-training model.

Examples

One health system noted that a 20–30 minute onboarding session is often sufficient if it includes live practice. Clinicians may learn to verbalize findings (e.g., "I see you're pointing to your left shoulder") or monologue these details after the visit to ensure the AI captures the necessary clinical data for complete documentation.

Another health system adopted a hybrid model combining self-start landing pages and tip sheets, weekly live sessions for those preferring structured onboarding, and post-adoption optimization support.

Note that a self-start pathway is only safe and scalable when the foundational technical work is already in place; additional optimization sessions may also be needed once end users have time to apply the software in their workflows.

BP36: Encounter-level verbal consent with written opt-out safeguards



Implement a consent framework that makes active, encounter-level verbal consent the standard practice whenever ambient AI is used. The provider should explain the technology's role, ensure the patient is


aware of its use, reaffirm the patient’s agreement on the recording itself, and document consent in the medical record. Written (or electronic) consent should be positioned as a complementary option, not a replacement for verbal consent. Several organizations are already utilizing statements about ambient AI in their boilerplate consent forms; however, the key is ensuring that patients are aware of this inclusion and that a reliable opt-out option is built into the consent process, including a method for opting out at any point in the future.

Note

Recording and transmission of data related to ambient AI can be presented as a separate checkbox in the standard form or obtained at check-in for each encounter. Written consent mechanisms are particularly important in cases where a patient is not able to provide verbal consent due to lack of capacity or a mental or physical condition affecting speech.

Note that health systems should recognize that consent requirements vary by state and by age of participant. Consent to treat vs. consent to record differs by age in some states.

BP37: Align consent/notification strategy with state law (one-party vs. two-party vs. mixed)



Align your consent and notification strategy with applicable state and local regulations before determining the appropriate level of transparency. In two-party consent states, passive notification is insufficient; health systems must implement affirmative confirmation of AI use such as combining physical signage in exam rooms with verbal healthcare provider confirmation and, where possible, visual indicators (such as an icon on a patient list) to signal to the entire care team that recording is active. In one-party consent states, passive notification may be legally adequate.

Note

Health systems with dedicated legal, regulatory, and privacy teams should leverage those resources to define and validate the consent strategy. Smaller or regional health systems that lack these internal capabilities should work closely with their ambient AI vendor for guidance on regional compliance requirements.

BP38: Flexible training framework reinforcing the human-in-the-loop paradigm



Establish a flexible training framework (e.g., multi-modal, tiered) that prioritizes asynchronous resources and on-demand, one-on-one expert consultation. Throughout training explicitly highlight that the

healthcare provider maintains ultimate professional and legal accountability for the accuracy, clinical validation, and integrity of the final medical record.

Example

To prevent implementation bottlenecks, one organization transitioned from lengthy mandatory instructional videos to a more flexible model featuring a centralized resource site and virtual "drop-in" rooms for real-time troubleshooting. Leadership consistently reinforced that while the AI streamlines documentation, the provider remains the legally responsible party for the final output, mirroring the professional responsibility required when supervising a human scribe.

BP39: Role-appropriate training, multiple support pathways, and peer community channels



Provide role-appropriate training (e.g., e-learning with an embedded attestation confirming understanding of consent, usage, and applicable local, legal, and regulatory requirements) and decide whether to position it as a prerequisite or strong recommendation based on risk, regulatory context, and the evolving maturity of clinician familiarity. Offer at least two distinct support pathways: (1) a self-guided model for confident, independent users, and (2) a higher-touch, at-the-elbow model for users needing more support. Organizations should also create and maintain active peer community channels to support ongoing knowledge-sharing, tips, and workflow optimization throughout the lifecycle.

Example

One health system built two distinct enterprise support models. The self-guided model provided e-learning modules, virtual office hours, and tip sheets. The at-the-elbow model layered in secure chat support and on-site Learning Services presence at high-volume clinician locations. All users were required to complete e-learning with an embedded attestation confirming they understood consent, usage, and applicable local, legal, and regulatory requirements before gaining system access. The organization also maintained a private peer community channel where staff shared tips, customization guidance, and workflow optimizations.

BP40: Embed communication in clinician channels and operationalize segmented dashboards



Consolidate ambient AI communication into the channels that clinicians already use to do their work. Target communications to specific user cohorts based on utilization data, and continuously evolve

monitoring dashboards so that the metrics being tracked translate directly into adoption, training, and support actions.

Example

One health system found that broad, asynchronous communications across multiple platforms were largely going unnoticed by clinicians. The organization shifted its primary communication channel into the EHR in-basket (the same surface where clinicians already triage their work) and used utilization data to tailor each message to the specific audience it was relevant to. In parallel, the organization evolved its dashboard from a generic utilization view to a multi-dimensional view that breaks usage down by setting of care (ambulatory, emergency department, inpatient), entry mode (native EHR mobile app vs. copy-paste), and clinician role. This view surfaced that inpatient functionality, although available for roughly six months, had drawn limited uptake. Another large academic medical center has built a real-time live data dashboard accessible to executives and to the implementation team, used not only to report progress but to identify low and high users and to drive proactive outreach.

Monitoring

BP41: Track locally relevant performance, wellness, and patient experience metrics



Responsible AI Lifecycle Management & Use

Track implementation success through a combination of metrics that are locally relevant, including accuracy of coding, error rate in medication orders, and other key metrics based on the feature set available to the health system.

Example

For example, one successful strategy involved using a "super user" team (e.g., an optimization team) to identify providers struggling with high documentation burdens and offering them targeted ambient AI support. Additionally, pushing automated, short surveys to patients 24 hours after an encounter helped the organization understand if the ambient technology is enhancing the patient-provider relationship.

BP42: Performance-drift detection paired with formal vendor communication strategy



Risk & Impact Assessments

Third Party Management

Implement a system (appropriate to the health system's available resources and capabilities) to track AI-generated documentation quality and utilization metrics in order to detect significant performance drifts or unplanned changes in the technology's behavior. In parallel, health systems should establish a formal communication strategy with their vendor to stay informed of product updates and their performance implications. When evaluating and onboarding vendors, confirm that they provide proactive release communications, including release notes with performance implications clearly documented, as well as training plans for new features or significant application updates to ensure the care team can adapt effectively.

Note

For health systems with advanced analytics capabilities, a "difference-in-differences" modeling framework can help identify flags that trigger a root-cause analysis, allowing the health system to determine if a change is a vendor-side update, a bug, or a shift in the clinical environment. For most health systems, however, a structured vendor communication cadence and vendor-provided release notes represent a more practical and accessible monitoring strategy.

Appendix 1: Supporting Resources

[U.S. State Consent Laws](#)

- *Tracker for all state laws around recording and consent. Updated as of June 2026.*

[Ambient AI Regulation Tracker](#)

- *Tracker for healthcare specific legislations. Updated as of June 2026.*

[AI Scribe Features](#)

- *This spreadsheet includes an institutional/user-focused evaluation template for healthcare delivery organizations. This spreadsheet was provided by the MedStar Health National Center for Human Factors in Healthcare. Updated as of June 2026.*

[Best Practice Statements – AI Lifecycle & Governance Domain Mapping](#)

- *This framework maps all 42 best practice statements in this Implementation Playbook to the AI lifecycle stage(s) they apply to and the CHAI AI Governance Framework domain(s) they are tagged against. Use this framework to navigate the playbook by lifecycle phase, by governance domain, or as a checklist for implementation planning.*

[Ambient AI Testing & Evaluation \(T&E\) Framework](#)

- *Includes a set of methods, metrics, measures, and/or benchmarks for developers and implementers to more concretely evaluate the responsible use of ambient AI technology.*

Appendix 2: AI Lifecycle

AI Lifecycle steps provide the structural backbone of this Implementation Playbook. Challenges/Insights and Best Practice Statements are listed at the earliest AI lifecycle step to which they apply. AI lifecycle definitions are included below:

Intake/Procurement

Not all AI solutions enter an organization through a formal procurement process; however, the initial step in managing any AI solution is a structured intake process. Intake consists of gathering and documenting essential information about the AI solution and its underlying model(s) to establish a baseline understanding of intended use, risk, and governance needs. CHAI recommends collecting this information through a model card (or comparable documentation) and, where applicable, standardized intake questionnaires.

Risk

Once intake information has been gathered, organizations should categorize the AI solution's risk prior to deployment. Risk categorization synthesizes information from the model card and intake questionnaires to assess risk domains. The resulting risk categorization informs which mitigation controls are required, the depth of evaluation and monitoring, and expectations across subsequent lifecycle stages. Additional guidance on risk categorization is detailed in CHAI's [Risk and Impact Assessment Governance Playbook](#) and made practical through CHAI's [Risk Categorization Tool](#).

Pre-deployment

Pre-deployment occurs after design, development, and assessment activities and before initiating a pilot. The goal of this stage is to assess readiness for real-world use at limited scale and to proactively address risks that may emerge once the AI solution interacts with production data, users, and workflows. Updates to clinical, population, and organizational risk should be incorporated based on insights from development, validation, and silent evaluation. This lifecycle stage applies to both developer and implementing organizations. It includes readiness assessments for the purchasing or deploying organization, clarifies roles and responsibilities between developers and implementers, and evaluates broader AI system design considerations such as safety, privacy, security, usability, and monitoring plans. For acquired or vendor-provided solutions, organizations may use this stage as part of procurement by requiring evidence of best practices and engaging jointly with developers to confirm appropriate planning, risk mitigation, and usability for the full AI system.

Additionally, evidence standards for validation should be clearly defined, including: (1) what constitutes sufficient validation evidence; (2) when local validation is mandatory versus accepted by vendor attestation; and (3) what documentation is required if local validation cannot be completed?

Piloting

Piloting is a controlled, limited-scale deployment designed to validate real-world performance, usability, workflow integration, and risk assumptions prior to full rollout. During this stage, organizations should

confirm that the AI solution is used as intended, evaluate early safety and performance signals, and assess human-AI interaction, including user reliance, overrides, and escalation behaviors. Pilot results should be reviewed against predefined success criteria, risk thresholds, and Responsible AI principles, and findings should inform decisions to proceed, iterate, expand, pause, or terminate deployment.

Pilot exit criteria should be predefined and include: measurable Responsible AI thresholds (e.g., fairness, safety, accuracy benchmarks), explicit go/no-go authority designations, and formal sign-off from the designated AI oversight body before proceeding to full deployment. Piloting is strongly encouraged in most cases and should be considered a requirement for high risk AI solutions prior to full deployment.

Deployment

Deployment represents the transition from pilot to broader operational use. At this stage, the AI solution is integrated into standard workflows, supported by finalized policies, training, and operational ownership. Deployment should be accompanied by activation of monitoring mechanisms, clear accountability for oversight and incident response, and communication to affected stakeholders regarding the AI's role, intended use, and limitations. Any material changes to scope, configuration, or intended use identified during deployment should trigger reassessment through intake or risk categorization processes as appropriate.

Monitoring

Monitoring is an ongoing lifecycle activity focused on evaluating safety performance, effectiveness, and responsible use over time. Organizations should regularly review defined metrics through dashboards or comparable reporting mechanisms, track usage patterns, and flag unexpected behavior or deviations from intended use for manual review by IT, compliance, and other oversight functions. Monitoring should account for model drift, workflow changes, user behavior, and downstream impacts, particularly where direct model telemetry is limited. Additionally, monitoring activities may include tracking performance over time, identifying potential bias or unintended consequences, and assessing continued alignment with clinical or operational workflows. Organizations may also define conditions that trigger re-review, modification, temporary suspension, or termination of AI tools.

Human factors monitoring (e.g., override behavior tracking and over-reliance/automation bias detection) should be incorporated into monitoring programs. Training updates should be triggered by monitoring findings that reveal problematic usage patterns.

Monitoring should vary in frequency and intensity based on the risk of the AI solution. Higher risk solutions require more frequent review cycles and enhanced escalation pathways. Risk-proportionate monitoring ensures resources are directed where patient safety exposure is greatest.

Termination (if applicable)

Termination involves the planned or unplanned retirement, removal, or disabling of an AI solution. Organizations should define thresholds and criteria for termination in advance, informed by safety concerns, performance degradation, changes in intended use, regulatory developments, or organizational priorities. Prior to sunset, organizations should assess the operational, clinical, and workflow impacts of removal and ensure appropriate transition planning, communication, and data

handling. Documented termination criteria, a stakeholder notification plan, archival of monitoring data, and a lessons-learned review are important components of any AI sunsetting process. Formal sunsetting processes and audit retention are industry best practice; AI removal can create clinical disruption and liability exposure if unmanaged.

Appendix 3: AI Governance Framework

CHAI's [AI Governance Framework](#) establishes 8 governance domains/subdomains and 40 baseline operational controls that describe what a mature AI governance program looks like inside a health system. The Best Practice Statements in this Implementation Playbook are tagged against the domains/subdomains in this governance framework so that Implementers can map operational guidance directly to their existing control environment. CHAI's AI Governance Framework domains/subdomains and associated baseline operational controls are included below:

Domain 1: AI Policy

- 1.1 Develop a formal AI Policy document, approved by organizational leadership, and make accessible to all relevant staff.
- 1.2 Define key concepts in the AI Policy (e.g., definition of AI, scope of AI solution use, AI approval process).
- 1.3 Establish an appropriate individual or individuals with designated oversight for policy changes.

Domain 2: Organizational Structures

- 2.1 Ensure your organization has appropriately trained staff assigned to critical roles and responsibilities (as defined by your organization). These can be existing or new staff, at the organization- or network-level, and can have one staff responsible for one-to-many roles/responsibilities.
- 2.2 Establish an AI Governance Committee – a formal committee structure bringing together all governance stakeholders and responsibilities as defined by the organization.
- 2.3 Establish an intake and monitoring process for thorough investigation of prospective solutions and mechanisms for incident detection, feedback, and follow-up.
- 2.4 Establish protocols for escalating and communicating concerns/incidents for resolution to the governance structure

Domain 3: Organizational Resources

- 3.1 Identify and document security infrastructure and methods (HITRUST, SOC 2, NIST CSF, CISA) available and needed to protect data; including methods for encrypting data at rest and in transit.
- 3.2 Identify and document the tooling and data resources available and needed to conduct regular security assessments and vulnerability scans to identify and address risks.
- 3.3 Identify and document tools, resources, and methods available and needed to implement AI cybersecurity evaluation (penetration testing, adversarial robustness, etc.).
- 3.4 Maintain comprehensive records of AI-related assets including datasets, models, and tools to ensure traceability.

- 3.5 Identify and document tools and resources available and needed for standard documentation of all AI systems and for integrating them into an AI inventory, such as adopting a Model Card.
- 3.6 Identify and document the tools and resources available and needed to implement and maintain an AI inventory registry to track internal and third-party models.
- 3.7 Identify and document the tools, resources, and infrastructure available and needed to carry out pre-implementation validation testing and monitoring for relevant AI models, including responsible AI baseline tests.
- 3.8 Identify and document the computing resources available and needed to conduct recurring cybersecurity audits of AI infrastructure; enforce access controls, and HIPAA compliance.
- 3.9 Establish an AI workforce competency framework (or related document) that defines: roles and responsibilities of AI governance personnel, minimum qualifications and/or certification. Include ongoing training programs to update personnel on governance concepts such as AI safety and bias mitigation.
- 3.10 Identify when clinicians and staff need to be trained on specific AI solutions prior to implementation; also, determine if certain AI solutions require regular training/competency assessments.

Subdomain 4.1: Responsible AI Lifecycle Management and Use

- 4.1.1 Develop and maintain a process that appropriately assesses and documents intended use of each AI system.
- 4.1.2 Define responsible AI principles, as aligned with the organization's mission and goal (e.g., usefulness, fairness, safety, transparency, security, privacy); document the objectives and processes that the organization will use to align with these principles across relevant lifecycle stages.
- 4.1.3 Operationalize a process to evaluate defined responsible AI principles (e.g., usefulness, fairness, safety, transparency, security, privacy) and reference industry recommended evaluation frameworks where available.
- 4.1.4 Define a process to monitor the safety, performance, and responsible use of AI-enabled solutions.
- 4.1.5 Establish a monitoring mechanism or process for performance reporting via a dashboard or similar interface. Monitoring mechanisms or processes should flag unexpected patterns for review by responsible parties.
- 4.1.6 Establish a process to inform relevant stakeholders about Responsible AI principles and processes (e.g., definitions, evaluation, monitoring).

Subdomain 4.2: Risk and Impact Assessments

- 4.2.1 Conduct a Risk Categorization assessment (pre-deployment) to categorize risk in the context of the organization's AI use case scope.
- 4.2.2 Document how Risk Categorization is evaluated, recorded, retained, and reviewed.
- 4.2.3 Leverage a rigorous Risk Assessment to evaluate AI solutions that are categorized as higher risk (as defined by the organization).
- 4.2.4 Document how Risk Assessment is evaluated, recorded, retained, and reviewed.
- 4.2.5 Conduct and document an AI System Impact Assessment that includes evaluation of risks and benefits for AI solutions.

Subdomain 4.3: Responsible Data Management and Use

- 4.3.1 Establish and maintain a structured data acquisition register tracking origin, handling, intended use, and stewardship of all data used in AI systems.
- 4.3.2 Establish comprehensive data use agreements (DUAs) defining guidelines and parameters for permissible uses, prohibiting re-identification and redistribution, requiring vendor security compliance, and including audit rights and data return/destruction terms.
- 4.3.3 Establish written HIPAA-compliant de-identification policies; evaluate PHI datasets to select the appropriate method; validate de-identification before AI training use; maintain continuous oversight of linked datasets.
- 4.3.4 Establish a process for recurring AI data quality and bias evaluation to help identify quality gaps and known sources of bias.
- 4.3.5 Establish a standard operating procedure (SOP) for AI data preparation including rules for data cleaning, encoding practices, normalization/standardization methods, and partitioning procedures

Subdomain 4.4: Third Party Management

- 4.4.1 Include clauses in all AI vendor contracts that specify clinical, operational and regulatory responsibilities; outline requirements for disclosure of known model limitations and risks; define liability for harm or failure resulting from the AI system.

Subdomain 4.5: Education, Training, and Feedback

- 4.5.1 Develop a process to ensure all relevant staff receive training on AI systems, including their purpose, limitations, and related information. Reinforce training via recurring briefings and optional seminars for deeper engagement.
- 4.5.2 Provide all patients, no later than the date of first service delivery, with a comprehensive Notice of Privacy Practices (NPP), including how protected health information (PHI) may be used by AI systems, in compliance with HIPAA.
- 4.5.3 Develop a process for patient transparency and/or consent where appropriate for relevant AI solutions.
- 4.5.4 Develop a process to document, escalate, and notify relevant parties about AI incidents. Ensure it is clearly communicated that all parties involved in an AI incident share responsibility for reporting any voluntary or mandatory incidents.
- 4.5.5 Ensure information included in AI incident communication falls under relevant federal and state laws and protections.
- 4.5.6 Provide whistleblower protections for both internal and external reporting, including anonymity, confidentiality, non-retaliation assurances, and clear disclosure of data retention policies.

Appendix 4: Consensus Method

Best practice statements are collected from work group presentations and discussions. To ensure alignment across stakeholders, CHAI uses a three-phase consensus process for Best Practice Statements (BPS) generated through work group activities:

Phase 1: Initial Consensus Check

- **Purpose:** Gauge initial agreement on each draft BPS.
- **Voting Options:**
- *Include / Include Contextually / Exclude / Abstain*
- **Decision Rules:**
 - If $\geq 2/3$ vote "Include" → **Consensus achieved** (no further action).
 - If $< 2/3$ "Include", but $\geq 2/3$ combined "Include" + "Include Contextually" → **Flagged for Phase 2.**
 - If $\geq 25\%$ vote "Exclude" → **Flagged for Phase 3.**
 - If $\geq 2/3$ vote "Exclude" → **Automatically excluded**

Phase 2: Revote with Revisions

- **Purpose:** Re-evaluate BPS that did not reach consensus in Phase 1 (but had $< 25\%$ "Exclude").
- **Format:** Original and revised BPS shown side-by-side (based on Phase 1 feedback).
- **Voting Options:** *Include / Exclude / Abstain*, with an optional comment field.
- **Outcome:** Results used to determine final inclusion or exclusion.

Phase 3: Live Discussion and Vote

- **Purpose:** Address BPS with $\geq 25\%$ "Exclude" in Phase 1 (strong disagreement).
- **Steps:**
 - Facilitated group discussion of flagged BPS.
 - Live revote during the meeting + optional 1-week offline voting.
- **Voting Options:** *Include / Exclude / Abstain*
- **Outcome:** Final decision made based on discussion and revote results.

Appendix 5: Thank You and Attribution

We want to start by thanking every individual who showed interest, participated, listened, and came along with us in our work. CHAI is at its core, a convener, and a member-driven non-profit. We are so grateful to be on this journey with you towards responsible AI in health for all. Your experiences, your feedback, your contributions, all make us who we are and help bring us to where we need to be.

CHAI extends its sincere appreciation to all participants for their time, expertise, and thoughtful feedback throughout the development process. CHAI would like to provide special recognition to the individuals below who requested attribution by name for their contributions to these efforts:

- Kang Hsu, MD, Canary Speech
- Meghan Reading Turchioe, PhD, MPH, RN, Columbia University School of Nursing
- Ashley M Hopkins, PhD, Flinders University
- Lucy Sutphen, MD, HealthPoint
- Craig Norquist, MD, HonorHealth
- Joe Derenzo, PMP, Healthcare Performance Group Inc.
- Renee Moss RN, BSN, CPhT, PMP, Healthcare Performance Group Inc.
- Jennifer Shannon MD, Individual Contributor
- Joan Chang, MS, MPhil, Individual Contributor

- Brittney Harrell, CISSP, Nabra
- Jeremy Attermann, MSW, National Council for Mental Wellbeing
- Jeffrey A. Gold, MD, Oregon Health & Science University
- R Logan Jones MD FACP, Oregon Health & Science University
- Taylor N. Anderson, MD, MS, Oregon Health & Science University
- Sudha Jayaraman MD MSc FACS, Suki AI
- Christopher H. Lee, MBA, BSN, RN-BC, UCLA Health
- Hwayoung Cho, PhD, RN, FAMIA, University of Florida College of Nursing
- Maritza M. Suarez, MD, University of Miami Miller School of Medicine
- Shiba Kuanar, University of Minnesota
- Ann Wieben, PhD, RN, NI-BC, FAMIA, University of Wisconsin-Madison

Appendix 6: Document Process & Governance

AI Transparency Statement: AI was used in the initial drafting, thematic analysis, and summarization to aid in the development of this document. AI was also used to summarize workgroup discussions and presentations into draft best practice statements. All best practice statements and initial drafts were reviewed and heavily edited by workgroup leads, members, and the CHAI program management team.

Document Maintenance Plan:

- BPG will be hosted on CHAI website.
- JIRA/Microsoft Issue Form will be available for ongoing feedback and new suggested content
- CHAI program management (PM) will monitor feedback submissions via internal tracker
- For any submission that is deemed “major”, CHAI PM will flag and raise with work group leads.
- CHAI PM and ad-hoc WG Leads will communicate via email with the flagged feedback submission.

CHAI PM will review all feedback and flag "major" feedback.

- For flagged feedback, 1-3 options will be considered after work group leads review:
 - No action needed
 - Straightforward feedback – can be resolved via email correspondence and CHAI PM reflects changes on website
 - Complex feedback – brief 30min meeting with WG Leads to discuss and resolve. CHAI PM reflects changes on website. This meeting can be used for 1:M complex feedback submissions. Formal decisions will require simple majority vote by WG Leads + CHAI PM