Companion to the AI Use Policy Template for Hospitals

Section-by-Section Guide

Includes Rural Proofing Additions

August 2026

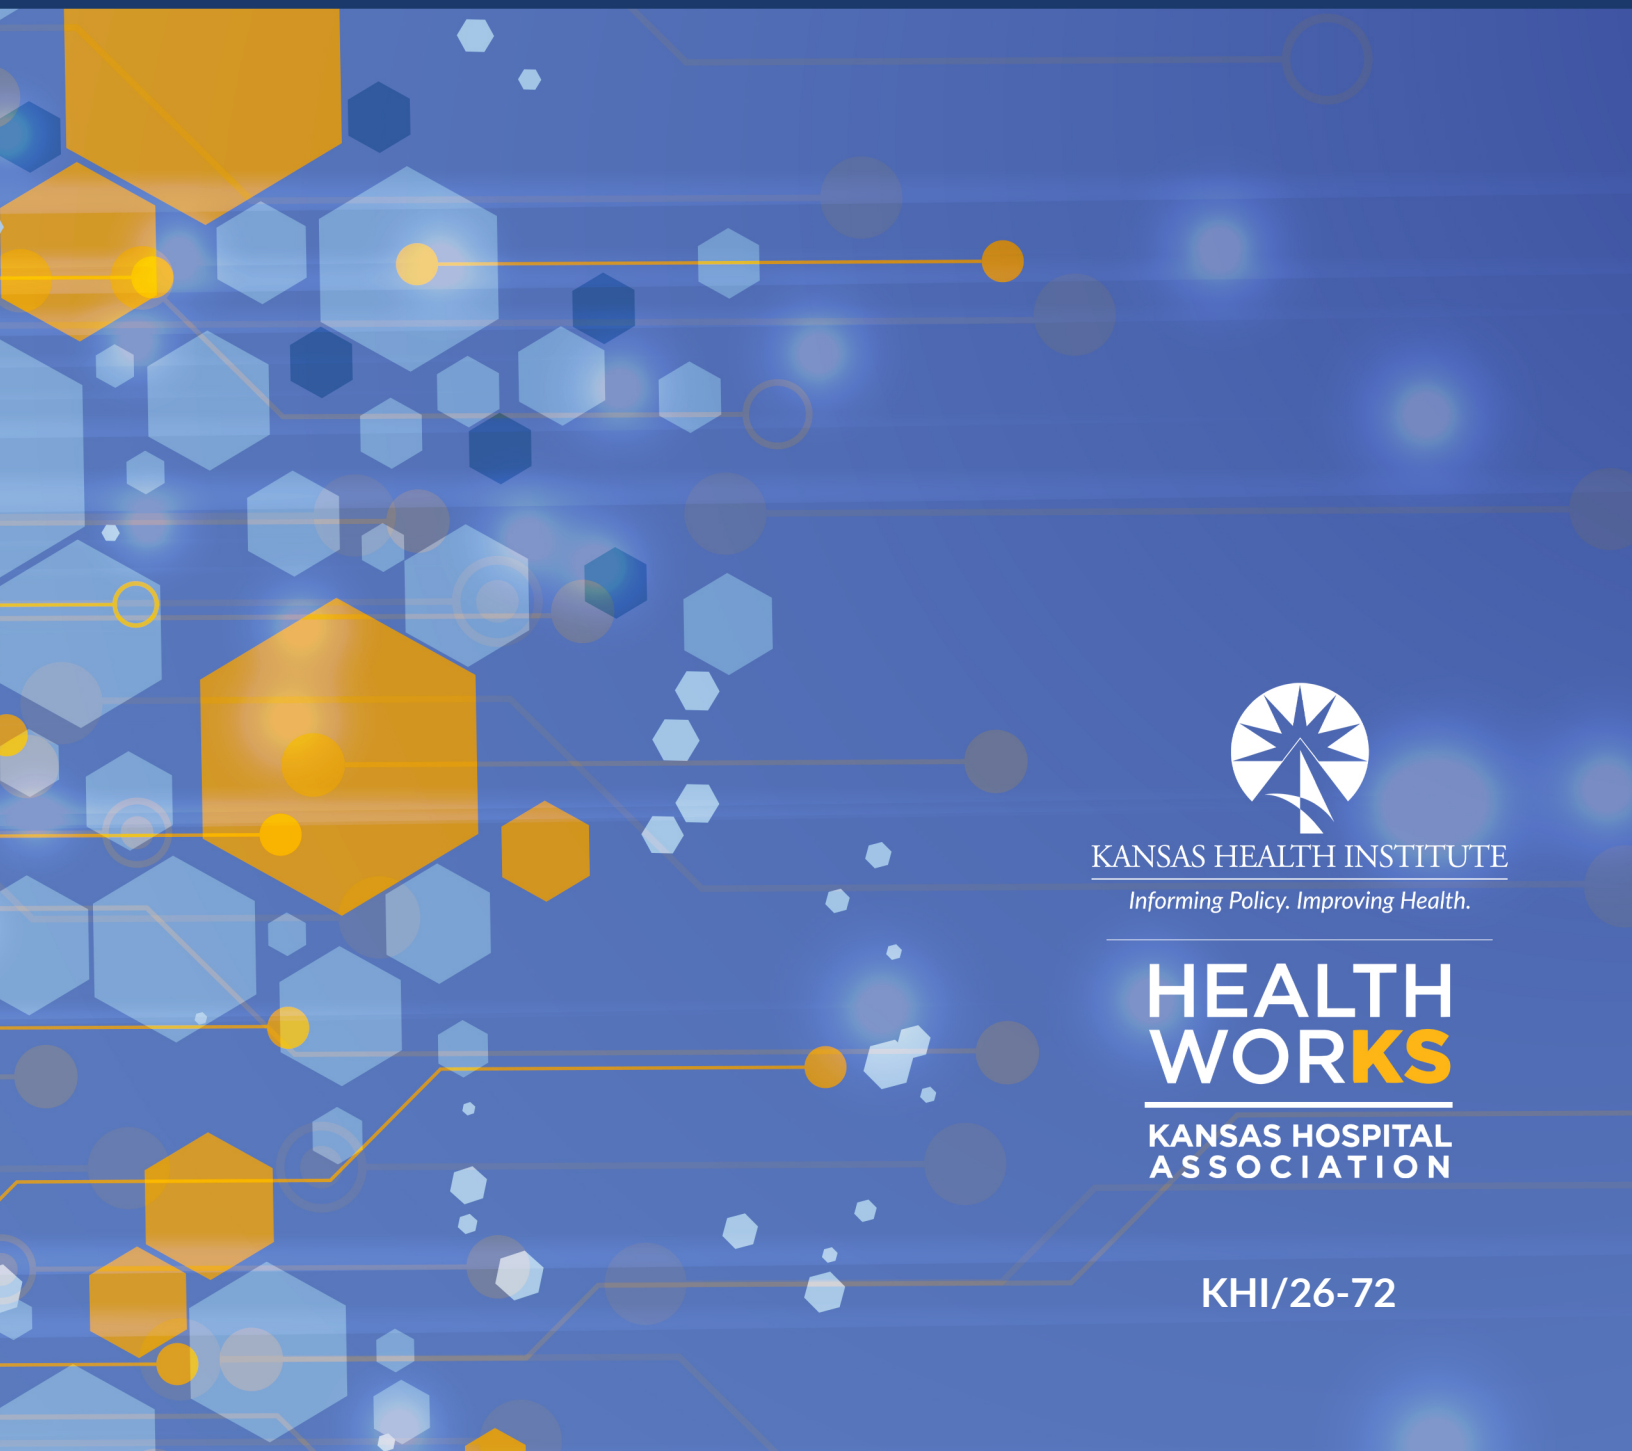

KANSAS HEALTH INSTITUTE

Informing Policy. Improving Health.

HEALTH WORKS

KANSAS HOSPITAL ASSOCIATION

KHI/26-72

Legal Disclaimer

This *Companion to the AI Use Policy Template for Hospitals: Section-by-Section Guide* is intended for educational purposes only and does not constitute legal advice. It is not a substitute for consultation with legal counsel or other qualified professionals. Hospitals should consult their own legal and other advisors before adopting or implementing an artificial intelligence use policy and should adapt this *Companion Guide* to meet the needs of their individual organization.

Because AI technologies, regulations and best practices continue to evolve, hospitals should regularly review and update their policies to ensure ongoing relevance and compliance. The authors make no representations or warranties of any kind, express or implied, regarding the content of this *Companion Guide*. Links to third-party content are provided for reference only, and the authors do not guarantee the accuracy or reliability of information provided by third parties.

About This Resource

The *Companion to the AI Use Policy Template for Hospitals: Section-by-Section Guide* (hereafter, the *Companion Guide*) provides section-by-section guidance for completing and implementing the *AI Use Policy Template*. It explains the purpose of each section, why it matters, and key considerations for selecting among the available options.

In this context, artificial intelligence (AI) refers to computer systems and technologies — ranging from predictive analytics to generative tools — that process data to perform tasks such as analyzing information, generating text or images, or supporting decisions and recommendations that would otherwise require human judgment.

The AI Use Policy Template for Hospitals: Core Policy + Modules (hereafter, the *AI Use Policy Template*) is a customizable policy document that hospitals complete and adapt for their organization. It is organized into clearly labeled sections with checkboxes and fillable fields to guide organizations through defining how AI will be used, who the policy applies to, and how to address allowable and prohibited uses using a risk-based framework. It is intended as a starting point and should be adapted to reflect your hospital's specific context, needs, governance structure, and applicable laws.

Throughout this *Companion Guide*, the term "hospital" is used because this resource was developed specifically for hospitals. However, the guidance is also applicable to health systems, which may adapt governance structures, responsibilities, and implementation approaches to reflect system-level operations. Together, the *AI Use Policy Template* and the *Companion Guide* are designed to support hospitals and health systems of varying sizes, capacities, organizational structures, and levels of AI experience.

This *Companion Guide* also addresses rural invisibility in AI, the risk that rural communities are underrepresented in the data, assumptions and governance structures that shape AI systems, making rural-specific needs and harms harder to detect. Throughout this guide, corresponding sections explain why rural context matters at each stage of the policy, from bias assessment and vendor procurement to patient disclosure, so hospitals can apply that reasoning to their own rural setting.

Authors

Tatiana Lin, M.A.
Director of Business Strategy and Innovation,
Kansas Health Institute

Emma Uridge, M.P.H.
Analyst,
Kansas Health Institute

Shelby C. Rowell, M.P.A.
Senior Analyst,
Kansas Health Institute

Acknowledgments

The *AI Use Policy Template* and the *Companion Guide* were developed by the Kansas Health Institute. Funding to support the development of these resources was provided by Healthworks.

The development of these resources built upon work previously completed by KHI through the Centers for Disease Control and Prevention (CDC) of the U.S. Department of Health and Human Services (HHS) Strengthening U.S. Public Health Infrastructure, Workforce, and Data Systems grant (OE22-2203). The contents are those of the author(s) and do not necessarily represent the official views of, nor an endorsement, by CDC/HHS, or the U.S. Government.

The Kansas Health Institute supports effective policymaking through nonpartisan research, education and engagement. KHI believes evidence-based information, objective analysis and civil dialogue enable policy leaders to be champions for a healthier Kansas. Established in 1995 with a multiyear grant from the Kansas Health Foundation, KHI is a nonprofit, nonpartisan educational organization based in Topeka.

Healthworks serves as the nonprofit arm of the Kansas Hospital Association. Healthworks educates partners, brings groups together and finds funding to provide services to change health care for the better. Healthworks believes with bold ideas and brave leadership, anything is possible.

This guide is available online at khi.org/articles/ai-use-policy-template-for-hospitals-and-companion-guide.

Who Should Use This *Companion Guide*

This *Companion Guide* is designed for the person or team responsible for drafting the AI use policy. That may be the Chief Compliance Officer (CCO), Information Technology (IT) Director, Chief Medical Officer (CMO) or a Compliance Analyst. Whoever leads the drafting should bring in the relevant stakeholders for their sections, particularly the CMO for clinical sections, the CCO for compliance sections and the IT director for the technical framework. In smaller hospitals and Critical Access Hospitals (CAHs), one or two people may cover all of these roles.

How to Use This *Companion Guide*

Work through each section of the *AI Use Policy Template* alongside the corresponding section of this *Companion Guide*. The *AI Use Policy Template* is organized as a series of bullet points with fillable checkboxes; check the boxes that apply to your hospital as you go. As you make decisions about which options to select, consider your hospital's needs, capacity, operational practices, and other relevant factors. Use your completed selections as the foundation for writing your policy.

Tailor the *AI Use Policy Template* to fit. Add, remove or modify sections and components within sections as needed to reflect your hospital's operations, responsibilities and risk environment. Most checklists in the *AI Use Policy Template* include an "Other: [ ]" option to allow hospital to add items that may not already be listed.

Plan for variation in completion time. Estimated times reflect a single reviewer working through each section and do not account for factors such as collecting and reviewing existing information, scheduling stakeholders, internal review cycles, or leadership approvals. All hospitals should treat these estimates as a starting point, not a project timeline, and plan accordingly based on their own size, staffing, and organizational complexity.

Before You Begin

Assemble your team before drafting. At minimum, the following roles should review relevant sections before the policy is finalized:

- Chief Compliance Officer (CCO) / Privacy Officer — policy coordination, Health Insurance Portability and Accountability Act (HIPAA) compliance, vendor obligations
- Chief Medical Officer (CMO) — clinical sections, patient disclosure, medical staff credentialing
- Information Technology (IT) Director / Chief Information Security Officer (CISO) — risk framework, security levels, AI Use Register, vendor technical review
- Chief Nursing Officer (CNO) — clinical documentation, alert fatigue
- Legal Counsel — contractual provisions, liability, Medical Staff Bylaw alignment
- Human Resources (HR) / Education — training requirements, staff acknowledgment

In CAHs and smaller hospitals, one person often holds multiple roles. That is fine — document clearly who is responsible for what.

Planning Tip

Break the drafting process into two or three working sessions rather than attempting to complete it in one sitting.

Suggested approach:

- Session 1 — Sections 1-9 (Purpose through Documentation Tools).
- Session 2 — Sections 10-19 (Risk Framework through Policy Review).
- Session 3 — Add-On Modules and Appendices.

Schedule legal and leadership review after Session 3 before finalizing.

About the Appendices

The *AI Use Policy Template* separates core policy content (Sections 1-19 and Add-On Modules A-E) from organization-specific information that changes frequently (Appendices A-F). When a governance role-holder changes, you should not need to initiate a full policy revision but focus on updating Appendix A. When a new AI tool is approved, you add it to Appendix B. This structure keeps the policy itself stable while allowing the operational details to stay current.

Figure 1. **Appendix Contents, Update Triggers and Responsible Parties**

Appendix	What It Contains	When to Update	Who Updates (Examples)
Appendix A: Governance Roles and Contacts	Names, titles and contact information for every governance role assigned in Section 7	When any named role-holder changes	Chief Compliance Officer
Appendix B: AI Use Register	Complete inventory of all approved AI tools with tier, Business Associate Agreement (BAA) status, oversight owner and approval date	When a tool is added, retired or reclassified	IT Director
Appendix C: Approved Patient Disclosure Scripts	CMO-approved verbal disclosure language for each ambient documentation tool	When a new ambient AI tool is approved or language changes	Chief Medical Officer
Appendix D: Regulatory References and Policy Cross-References	Hospital-specific policy numbers mapped to regulatory citations	When internal policy numbers change	Chief Compliance Officer
Appendix E: Approved Tool Lists by Module	Module-level tool inventories for each selected Add-On Module (A through E)	When tools are added, retired or changed	IT Director
Appendix F: Control Role Assignments	Responsible role for each required Tier A and Tier B control in Section 10	When roles or personnel change	Chief Compliance Officer

Each appendix includes an 'Appendix Version Date' and 'Last Revised By' field. Update both every time the appendix changes. Communicate every appendix update to the policy owner (typically the CCO) within five (5) business days of any change.

Implementation Phasing

As you develop your AI policy, you may consider a tiered approach in which core policy components are established first, while more detailed operational processes, procedures and supporting tools are phased in over time. Additional components may include vendor assessment checklists, the intake process and AI Use Register, staff training materials, incident reporting procedures and public communications about AI use.

Figure 2. **Example Phased Approach to AI Governance Implementation**

Phase	Key Actions	Timeline
Foundation	Designate governance roles; complete risk framework; establish prohibited use rules; conduct initial staff training; notify staff of the policy; file all staff acknowledgments, complete Appendix B (AI Use Register) with all active tools	Months 1-3
Operational	Notify vendors of transition period; incorporate AI contract clause into new contracts; communicate publicly how the hospital uses AI	Months 4-12
Continuous Improvement	Annual policy review; refresh training materials; reclassify tools as scope changes; incorporate Board AI report; document lessons learned from deployed systems	Year 2+

Section-by-Section Guide

Header: Policy Header and Administration

Estimated Time 15-30 minutes

The policy header table captures administrative information that tracks the policy's status over time. Complete every field.

Policy Number

Assign a unique identifier. A simple format such as AI-2026-001 makes it easy to reference in training materials, incident reports and other documents. If your hospital already has a policy numbering convention, use that.

Version

Start at 1.0. Use minor version numbers (1.1, 1.2) for small clarifications and major version numbers (2.0) for significant revisions. The Record of Revisions at the end of the *AI Use Policy Template* is your audit log in order to keep it current.

Approving Authority

This should be the highest-ranking official accountable for this policy such as the CEO or Board of Trustees Chair. Having a named approving authority establishes that the policy carries organizational weight.

Next Review Date

Annual review is the standard recommendation for AI policy given the pace of change. Set the Next Review Date when the policy is approved and build it into your compliance calendar immediately.

Section 1: Purpose

Estimated Time 15-20 minutes

The purpose section, including the policy purpose statement, establishes the intent and authority of the AI policy, clarifies why it exists and sets the overarching direction.

Hospital Legal Name

Enter the full legal name of the hospital in the blank at the top of Section 1. Use the full legal name as it appears in your licensure documents and Joint Commission credentials.

Commitment Statement Check Boxes

Select all statements that would guide your hospital in implementing and governing AI use across the organization. At minimum, we recommend selecting Patient Safety First, HIPAA Compliance, Human Oversight and Transparency.

Tip: Align With Existing Mission Language

If your hospital has adopted specific equity, patient safety or community trust language in other policies or your strategic plan, use consistent language here.

Section 2: Scope

Estimated Time 20-30 minutes

The Scope section defines what AI technologies this policy governs. AI is increasingly embedded in tools hospitals already use, sometimes without the vendor labeling it as AI. Sepsis prediction models built into your electronic health record system, natural language processing (NLP) in your documentation platform, and predictive scheduling tools are all AI even if they are not marketed that way.

In the provided list select all technology types that should fall within scope. Note that the purpose is to identify categories, not list specific tools and systems.

EHR-embedded AI note: Tools like Epic's Sepsis Prediction Model or Cerner's AI modules are AI tools under this policy even though they are licensed as part of your EHR contract. They must appear in Appendix B (AI Use Register) and be assigned a risk tier. Your IT Director or EHR administrator can help identify which AI modules are active in your system.

Section 3: Applicability

Estimated Time 15-20 minutes

This section defines who is governed by the policy. Coverage should extend to anyone who uses, interacts with or handles outputs or data from AI tools and systems.

3.1 Departments and Care Areas

Check all departments and care areas to which this policy should apply. If your hospital has specialty departments not listed (e.g., behavioral health, rehabilitation, long-term care units), add them under "Other." If the policy is intended to apply hospital-wide, this section may not be necessary. However, this section can help clarify scope when the policy applies only to specific departments or care areas.

3.2 Workforce Coverage

In determining the groups identified in Section 3.2, hospitals may consider including employees, contractors, consultants, volunteers, interns, fellows, trainees and other individuals or organizations performing work for or on behalf of the hospital who may use AI tools or interact with AI-generated outputs as part of their responsibilities. This can help promote consistent expectations and responsible use practices.

Particular attention is warranted for three groups: (1) traveling and agency staff, who may not receive your standard onboarding; (2) students, residents and trainees, who are supervised but may independently use tools; and (3) contractors and consultants, who may bring their own AI tools into your environment.

3.3 Vendors and Third Parties

This section notes that vendors are not directly bound by this policy but are governed through procurement requirements, contract terms and vendor assessments. Where AI is embedded in an existing platform, departments may have limited ability to negotiate vendor terms or data practices — in those cases, hospitals should review vendor disclosures to determine whether they are acceptable or whether an alternative vendor should be considered.

Section 4: Alignment with Existing Policies and Regulations

Estimated Time 15-20 minutes

This section anchors your AI policy within the existing regulatory and policy framework governing your hospital. Check all that apply. For most hospitals, this will include HIPAA, Joint Commission, CMS Conditions of Participation and applicable state-specific hospital licensure regulations.

Check for State-Level AI Guidance

For example, in Kansas, consult the Kansas Hospital Association, the Kansas Department of Health and Environment and your hospital's legal counsel to identify any applicable state-level AI requirements before finalizing your policy.

Section 5: Guiding Principles

Estimated Time 20-30 minutes

The guiding principles communicate the values shaping how your hospital uses AI. They are not just aspirational statements. Each principle has practical implications that run through the risk framework, training requirements and oversight obligations that follow.

Select the principles that genuinely reflect your hospital's mission and organizational commitments. You may select all of them. If your hospital has existing mission or values language that maps to one of these principles, incorporate it using the 'Other' field.

Four principles warrant emphasis in the hospital context:

- Patient Safety First: This is the foundational principle for hospital AI governance. Every AI safeguard in this policy traces back to it.
- Human Oversight and Clinical Accountability: Clinicians retain responsibility for every patient care decision, regardless of what AI suggests.
- Bias Awareness and Mitigation: Hospitals that serve diverse populations should pay particular attention to this principle. AI systems trained on non-representative data can produce systematically biased outputs that affect clinical decisions. Section 15.2 and Module B address this operationally.
- Workforce Sustainability: Staff concerns about AI replacing clinical roles are common. Naming this principle explicitly in the policy addresses those concerns directly and helps with buy-in during rollout.

Rural and low-volume populations are frequently underrepresented in the data used to develop and validate AI tools, a gap sometimes called "rural invisibility." A tool that performs well in a large urban health system's data may perform differently, unpredictably or less accurately in a smaller or more rural patient population. Hospitals serving populations that may be underrepresented in AI training or validation data, including rural communities, should treat this as a distinct dimension of bias review rather than assuming general bias-mitigation language already covers these risks. Hospitals should ask vendors directly whether and how a tool was validated on rural, low-volume or critical-access populations and document the answer even when it is "not validated."

Section 6: Key Definitions

Estimated Time 20-30 minutes

Shared vocabulary reduces errors. When every staff member means the same thing when they refer to "Procured AI Tool" or "Human-in-the-Loop," the likelihood of miscommunication and missed safeguards decreases. The definitions in the *AI Use Policy Template* use standard terminology drawn from HIPAA, the National Institute of Standards and Technology (NIST) and industry sources.

If your hospital uses different terms for any defined concept, revise the definitions to match your local usage, or add your terms to the additional definitions table at the bottom of the section. The table is there precisely for hospital-specific terms, EHR module names or state law definitions not captured in the standard list.

Three Definitions Deserve Special Attention

The following three terms deserve particular attention because the value of distinguishing between them in AI policy may not always be immediately intuitive. Without clarifying the differences between **publicly accessible AI tools and systems**, **formally procured AI tools and systems**, and **personal paid access to publicly accessible AI tools and systems**, staff may incorrectly assume the same approvals, safeguards, oversight expectations or data protection requirements apply equally across all types.

Section 7: Governance

Estimated Time 30-45 minutes

Section 7 establishes who is responsible for what is in your hospital's AI governance. The goal is a structure that is both clear and complete — one where every core function is assigned and people understand their responsibilities.

The *AI Use Policy Template* does not prescribe specific role titles because hospitals have different organizational structures, staffing levels and terminology. Each core governance function should reasonably be covered within the scope of the policy.

Completing Appendix A — Governance Roles and Contacts

The governance role assignments, including delegated authority, are included in Appendix A so they can be updated when personnel change without requiring a full policy revision. When completing Section 7, note the role titles in the policy itself. Go to Appendix A to assign specific names, direct contact information and backup/delegate designations.

Leveraging Existing Community Advisory Structures

For hospitals with an existing community advisory structure, such as a hospital association-affiliated workgroup or a hospital community benefit committee, consider routing high-impact AI tool decisions (Tier A deployments in particular) through those channels for advisory input before finalizing approval. This input is intended to inform, not replace, the hospital's existing approval authority under Section 7. Where a time-sensitive Tier A decision cannot reasonably wait for the advisory body's next scheduled meeting, the hospital may proceed under standard governance approval and provide the advisory body a summary of the decision at its next meeting for retrospective input.

CAH and Small Hospital Note

In Critical Access Hospitals and smaller community hospitals, one person often holds multiple governance roles. Clear documentation of these responsibilities is important. A potential structure for a small hospital:

- CEO, Administrator, or Executive Leader: Final AI policy authority; approval of Tier A tools
- Compliance, Quality, Privacy, or Risk Management Lead: Day-to-day AI coordination, AI Use Register oversight, BAA reviews, and policy implementation
- IT Director, IT Manager, or Shared Services IT Team: Technical review, cybersecurity oversight, vendor security assessment, and incident response
- Clinical Leadership (e.g., Chief Medical Officer, Medical Director, Chief Nursing Officer, Nursing Director, or Quality Leader): Clinical AI oversight, patient safety monitoring, and review of Tier A clinical AI tools
- Human Resources, Education, or Workforce Development Lead: Training completion tracking and staff acknowledgments

Where one person covers multiple roles, identify a backup for critical functions so governance responsibilities can continue during staff absences.

Rural and Critical Access Hospitals face a second, distinct risk beyond governance capacity. Many AI tools are validated primarily on data from larger, higher-volume health systems. When evaluating a tool, ask the vendor directly whether it has been tested on populations and care settings comparable to yours and document the answer, even where the honest answer is "not tested." Many CAHs lack the resources to run independent local validation studies, which makes vendor disclosure at procurement especially important.

7.2 Enforcement and Corrective Action

Select the enforcement mechanisms that are consistent with your hospital's existing HR and personnel policies.

Every enforcement mechanism you select must be consistent with your Medical Staff Bylaws, employment agreements and applicable state employment law. If your Bylaws do not currently address AI-related peer review, add it to your next Medical Staff Bylaws review cycle.

Section 8: Acceptable and Prohibited Use Overview

Estimated Time 15-20 minutes

Section 8 is the policy's quick-reference summary of what AI may and may not be used for. It is not intended to be exhaustive. The detailed risk classification rules are in Section 10 and the detailed prohibited use rules are in Section 13.

8.1 Permitted Use Check Boxes

Select all conditions that must be met for AI use to be permitted. At minimum: assigned risk tier with appropriate controls, completed BAA where protected health information (PHI) is involved, human oversight proportional to risk, and tool listed in AI Use Register.

8.2 Prohibited Without Approval

Under prohibited categories, select activities or uses for which AI is not permitted and add others as appropriate for your department.

Section 9: AI Documentation and Notetaking Tools

Estimated Time 20-30 minutes

Section 9 addresses two distinct AI documentation contexts: clinical ambient scribes (tools that generate documentation entered into the EHR) and meeting/administrative notetakers (tools that record and transcribe workplace meetings). Both create risks that are easy to overlook.

9.1 Clinical AI Documentation (Ambient Scribes)

List all approved ambient documentation tools in the table in Appendix E (e.g., Nuance DAX, Abridge, Suki). If no tool has been formally approved yet, write "None approved — pending CCO and CMO review." Do not leave the field blank. A blank field suggests the question was not considered.

The requirements checklist addresses the key clinical documentation compliance obligations. Check all that apply. At minimum, all of the following should be checked for any hospital using ambient AI:

- BAA coverage for the documentation tool
- Clinician review and attestation before filing in EHR
- Patient disclosure before the tool is activated
- Mechanism to document patient declination

Patient disclosure for ambient documentation tools: Approved disclosure scripts for each authorized tool are maintained in Appendix C. No ambient tool may be used without an approved script in Appendix C. Clinical staff may not improvise their own disclosure language — use only the CMO-approved script for each tool.

9.2 Meeting and Administrative AI Notetakers

AI notetaking tools have become common in workplace meetings. This section addresses them specifically because they create a privacy risk that is easy to overlook — they capture spoken conversations, which often include sensitive information that staff would never deliberately type into a system but may discuss freely in a meeting.

The approved script in Section 9.2 of the *AI Use Policy Template* is a practical tool for staff who feel uncertain declining an external partner's request to use their AI notetaker. Having an approved, professionally worded statement removes ambiguity and gives staff clear language to use without having to improvise.

Staff using AI notetakers should carefully consider what information may be discussed during the meeting and ensure the security, privacy and access controls of the notetaking tool align with the sensitivity of the information.

A simple default for meeting leads: State in the calendar invitation whether an AI notetaker will or will not be used, for any meeting involving external participants. Do not leave it unstated. This prevents last-minute confusion, ensures all participants have advance notice, and gives anyone with an objection the opportunity to raise it before the meeting begins.

Section 10: Risk and Security Framework

Estimated Time 45-60 minutes

The Risk and Security Framework is the operational foundation of the policy. Its core principle is that higher-risk tasks require more secure AI systems and greater oversight.

Note: Because the Risk and Security Framework establishes how AI uses are classified, governed and authorized, organizations should strongly consider involving legal counsel in the review and customization of this section to ensure alignment with applicable laws, regulations, contractual obligations and organizational risk tolerance.

Understanding the Two-Axis Framework

The framework works on two axes. Task Risk Tiers (Sections 10.1–10.3) describe what the AI is being used for, including the sensitivity of the data involved and the potential consequences if the output is wrong. Tool Security Levels (Section 10.5) describe the minimum security requirements a tool must meet before it can be used for work at that risk level. Rather than describing categories of tools, they establish a pass/fail standard that any tool must satisfy before it can be used for work at that risk level.

Figure 3. **Task Risk Tiers and Tool Security Levels for Hospital AI Use**

Tier	Applies To	Hospital Examples
Tier A High Risk	Direct patient care, clinical decision-making, PHI processing, functions affecting patient safety or legally protected rights	Sepsis prediction models, radiology AI, ambient documentation scribes, prior authorization AI, AI-assisted coding, AI-generated diagnoses or treatment recommendations
Tier B Moderate Risk	Core operations, public-facing content, revenue cycle, scheduling — where outputs shape decisions but do not directly affect patient care	Patient scheduling AI, staff scheduling models, AI-assisted communications drafts
Tier C Low Risk	Internal and non-sensitive drafting, brainstorming, summarizing public documents — no PHI, no personally identifiable information (PII), no official use without human review	Drafting internal meeting agendas, brainstorming staff newsletter topics, grammar checking non-sensitive internal documents, summarizing publicly available clinical guidelines

When in Doubt, Classify Higher

If a task falls on the boundary between Tier A and Tier B, classify it as Tier A. You can reclassify downward after review if appropriate. Under-classifying a task that involves health data is a significantly more serious error than being overly cautious.

Rural and Low-Volume Population Note

A tool's risk tier not only depends on what it does but also on whether it has been shown to work reliably for the population it will actually serve. If a vendor promoting a Tier A tool cannot provide performance data for rural, low-volume or critical-access settings, treat that absence of data as a documented limitation. Lower historical service use among underserved communities may reflect access barriers rather than lower clinical need; do not let sparse local data be read as low risk.

10.5 Tool Security Levels

Responsibility for evaluating and classifying the security, privacy and organizational risk level of AI tools should reside with the hospital role or governance structure designated to oversee AI, technology, privacy or information security governance — not with individual staff. One important distinction that many hospitals may not initially recognize is:

> A personal paid subscription to an AI tool (e.g., a clinician's own ChatGPT Plus or Claude Pro account) is a Level 3 tool — even though the employee is paying for it. The determining factor is not cost but whether the tool was formally procured by the hospital with a contract that includes data protection requirements and BAA coverage. If it was not procured by the hospital, it is Level 3 and may only be used for Tier C tasks with no PHI or PII.

Figure 4. **Task Risk Tier and Tool Security Level Authorization Framework**

Task Risk	Tool/System Security Level	Permitted?	Reason
Tier A (high risk)	Level 1 only	Allowed	High-risk tasks require the highest-security tools and systems
Tier A (high risk)	Level 2 or 3	Prohibited	Lower-security tools and systems lack the safeguards required for sensitive data
Tier B (modate risk)	Level 1 or 2	Allowed	Moderate-risk tasks may use procured tools and systems with confirmed privacy practices
Tier B (modate risk)	Level 3	Prohibited	Public or consumer tools and systems may not be used for tasks that touch operational or public-facing content, though exceptions are possible.
Tier C (low	Any level	Allowed	Low-risk tasks may use any tool and system type

If your hospital has operational reasons to be more restrictive, adjust accordingly. What matters is that you set the matrix deliberately and enforce it consistently.

Section 11: Clinical Competency and Acceptable AI Use

Estimated Time 20-30 minutes

Section 11 establishes that AI is a supplement to clinical competency, not a substitute for it. A clinician who relies on an AI tool to perform a task they cannot independently evaluate may be unable to provide the level of meaningful human oversight required by the policy.

11.1 Clinical Staff Requirements

Check all requirements that apply. At minimum: completion of required training before any AI tool use, and use of only approved tools listed in the AI Use Register (Appendix B).

11.2 Restrictions for High-Risk Clinical Tasks

The restrictions in 11.2 address the highest-risk intersection of AI and clinical practice. At minimum: AI may not be used in place of foundational clinical competency or independent professional judgment, and supervising physicians remain fully responsible for AI-assisted outputs produced by trainees under their supervision.

Practical Guidance for Supervisors

When onboarding new residents, fellows or clinical staff, introduce Tier A AI tools only after the individual has demonstrated baseline competency in the relevant clinical task independently. For example, a resident should be able to interpret a basic chest X-ray without AI assistance before relying on radiology AI to support that work. AI is a check on competent clinical judgment — it cannot replace the judgment itself.

Section 12: Training and Readiness Requirements

Estimated Time 30-45 minutes

Section 12 establishes that staff must complete training before using any approved AI tool, and that completion must be documented.

12.1 Required Training Topics

The four required training components are intentionally broad to give your hospital flexibility in delivery. For most hospitals:

- AI literacy and responsible use — available from many free sources including the American Hospital Association (AHA), the American College of Healthcare Executives (ACHE) and the American Medical Informatics Association
- HIPAA and data privacy in AI-enabled environments — can be integrated into existing annual HIPAA training
- Hospital AI policy and acceptable use — must be your hospital-specific policy content
- Tool-specific clinical training — must address each Tier A tool individually before first use

General AI literacy training covers how AI tools work and their limitations broadly. For hospitals serving rural communities, consider a distinct training component addressing how AI-influenced decisions intersect with rural-specific access barriers. If a no-show prediction model is built without transportation-access or travel-time data, it can't distinguish a patient who chose not to come from one who couldn't physically get there. A predictive risk score might flag a missed follow-up as "non-adherence" when the real driver is distance to care. Catching that distinction is a different skill than general AI literacy, and it's worth naming separately in training materials.

12.2 Training Completion and Acknowledgment

Training completion does not require a formal certificate. Staff can sign or electronically acknowledge training completion. The Staff Acknowledgment Form at the end of the *AI Use Policy Template* can be used for this purpose. A Learning Management System record or an email confirmation could also satisfy this requirement. What matters is that the record is stored centrally and can be produced in a survey or audit.

12.3 Tier A Tools — Additional Requirement

Before using any Tier A AI tool, staff must receive written authorization from the CMO or designee in addition to completing training. Fill in the blank with the title of the approving authority. This step ensures a clinical supervisor has confirmed that the staff member is ready and that appropriate tools and safeguards are in place.

Do not allow staff to begin Tier A AI-supported clinical work before both training acknowledgment and written authorization are on file.

12.4 Refresher Training Triggers

Check all circumstances that will trigger refresher training. At minimum: new Tier A tool or system deployment, major update to an existing approved tool, and AI-related incident. Annual policy review should also be included as a standard trigger.

Section 13: Misuse of AI

Estimated Time 20-30 minutes

13.1 Prohibited Personnel Conduct

Section 13.1 covers prohibited personnel conduct — specific behaviors that are not permitted regardless of the AI tool involved. This section applies to all Covered Persons as defined in Section 3.2.

Some of the most significant risks associated with AI use in hospitals involve situations where AI-generated outputs directly influence clinical decision making without adequate human oversight or validation. For example, risks may arise when:

- **AI-generated clinical recommendations are inaccurate, biased or hallucinated but are accepted without adequate human review**
- **AI-assisted diagnostic interpretation errors occur in areas such as radiology, pathology or electrocardiogram (ECG) interpretation**
- **AI-generated medication dosing or treatment recommendations contribute to patient harm.**

These risks may result in patient safety concerns, inappropriate treatment decisions, delayed care or adverse clinical outcomes if appropriate governance, validation and human oversight processes are not in place.

The prohibited conduct list in Section 13.1 covers real patterns of misuse observed in health care settings. A few points warrant emphasis during rollout:

- **Bypassing procurement or IT review:** Shadow AI, such as staff using tools that have not gone through the hospital's review, security or compliance process, is one of the most common governance gaps. The policy's procurement and approval requirements (Section 16) exist specifically to reduce security, privacy, operational and patient safety risks associated with unapproved

AI systems. Shadow AI is often difficult to detect through policy alone. Organizations should work with IT, compliance, privacy and security teams to determine whether existing monitoring, software inventory, network management, procurement or audit processes can be leveraged to identify the use of unapproved AI tools.

- **AI-generated misinformation and unvalidated clinical content:** Hospitals bear a unique public trust responsibility. AI-generated clinical recommendations, patient education materials, discharge instructions, public communications or other health-related content that are inaccurate, biased, incomplete or hallucinated may contribute to patient harm, misinformation or inappropriate clinical decisions if not properly reviewed. All patient-facing and clinically relevant AI-assisted content must undergo appropriate human review and validation prior to use or distribution.
- **Harassment, discrimination and biased outputs:** AI tools may generate discriminatory, biased, harassing or otherwise inappropriate content involving patients, staff or communities. This includes discriminatory employment-related uses, inappropriate patient-facing content or the generation of offensive text, images or recommendations. Such conduct constitutes serious misconduct under this policy.

The no-retaliation language at the end of this section is important. Staff who report suspected violations or concerns in good faith must be protected from retaliation. Without that protection, incidents may go unreported and organizational risks may remain unidentified or unmanaged.

Section 14: Vendor and Third-Party Obligations

Estimated Time 30-45 minutes

This section extends the policy's requirements to the consultants, contractors and vendors who work on the hospital's behalf. Many hospitals work with consultants, contractors, vendors and other third parties to support clinical, operational, administrative, technology, billing, revenue cycle, data management and patient care functions. The following are example definitions. Hospitals should align them with existing internal definitions, procurement standards and organizational policies where applicable.

- Consultants — Individuals or organizations that provide specialized professional, clinical, technical, strategic, legal, compliance or advisory services to the hospital, typically under a formal agreement or scope of work
- Contractors — Individuals or entities engaged by the hospital through a contract to perform specific clinical, operational, administrative, support or technical services on behalf of the hospital
- Vendors — External companies, providers or suppliers that sell, license, implement, support or maintain products, medical technologies, software, AI systems, platforms, equipment or services used by the hospital

14.1 BAA and Contractual Requirements

Check all requirements that apply. At minimum: valid executed BAA before any PHI is shared or processed, and prohibition on using hospital data to train AI models without written authorization. The data training prohibition deserves particular attention. Many AI vendors include data training provisions in their standard terms of service. Review all vendor contracts for this language and negotiate its removal or restriction before finalizing any agreement.

14.2 Required Contract Language

The sample contract clause in Section 14.2 is a starting point. Work with your procurement or legal office to incorporate it as standard language in all new vendor contracts.

14.3 Transition Period for Existing Contracts

In this section, the focus is on setting a transition period for existing contractors and consultants. This is the window of time after the policy's effective date during which vendors with active contracts must come into compliance. Timeframes for contractors and consultants to align with updated AI requirements may vary depending on the size, scope and complexity of the services or agreements involved. Fill in both blanks in Section 14.3 with the periods your department will use and fill in the title of the person who will review extension requests.

Practical Vendor Management Tip

Create a one-page AI Vendor Disclosure form your procurement staff can use with every existing vendor relationship. Key questions to ask:

- Does your product or service currently use AI or machine learning?
- What hospital data do you retain, and for how long?
- Do you use hospital data to train or improve your AI models?
- What is your breach notification timeline?
- Do you offer a no-train option to prevent our data from improving your models?

Send this form to all active vendors within the first 60 days after policy adoption.

Section 15: Responsible Use, Safeguards and Transparency

Estimated Time 30-45 minutes

Section 15 provides the operational safeguard requirements that support the framework established in Sections 7-14. It addresses human oversight, bias mitigation, transparency to patients and staff, and documentation of AI-assisted decisions.

15.1 Human Oversight Requirements

Select all oversight requirements that apply. At minimum: all AI systems operate under meaningful human oversight; clinicians review and validate AI-assisted findings before they affect patient care; and Tier A tools require documented review before deployment.

Fill in the escalation contact for AI errors or safety concerns. This should be a named role (e.g., "charge nurse, then CMO") — not just a department. Staff need to know exactly who to call when an AI tool produces an output that appears wrong or inconsistent with clinical judgment.

15.2 Bias Mitigation

Proactive bias assessment is required before deploying any Tier A tool. Before purchasing, ask vendors to disclose known limitations and populations where the tool's performance may be reduced. Document the results of this review in the AI Use Register (Appendix B). After deployment, the CMO and CNO should monitor outcomes for disparate patterns at least semi-annually and report findings. This review should include geography and service area as a bias dimension, not only demographic categories, particularly for hospitals where a meaningful share of patients are served under rural or critical-access conditions.

15.3 Transparency to Patients and Staff

Check all transparency commitments that apply. At minimum: patients are informed when AI materially influences their care, and the hospital maintains an AI Use Register (Appendix B) accessible to

authorized staff. The public-facing summary of hospital AI use is good practice and should be posted on your hospital's website.

15.4 Tier A Documentation Requirements

For each Tier A AI application, hospitals should maintain internal documentation identifying the tool name, intended purpose, level of risk, PHI involvement, applicable safeguards and responsible oversight authority. This documentation supports organizational governance, auditing, compliance and risk management activities.

Where required by hospital policy, regulation, accreditation standards or clinical workflow, documentation may also indicate when AI materially informed high-risk clinical decision making. Documentation requirements should be proportionate to the level of clinical, operational or patient safety risk associated with the AI application.

Example documentation language for selected high-risk clinical use cases may include:

- "AI was used to support [describe task] by [describe AI role]. Final clinical decisions were made by the responsible licensed clinician at [Hospital Name]."
- Hospitals should determine which AI use cases require formal medical record documentation versus internal governance documentation to avoid unnecessary documentation burden while maintaining appropriate accountability and human oversight.

Section 16: Acquisition of AI Tools and Systems

Estimated Time 30-45 minutes

Section 16 establishes that acquiring AI tools is not a decision individual staff or department heads make on their own. It requires involvement from technical reviewers, governance leads and procurement staff, with the level of review proportionate to the risk tier.

16.1 Pre-Procurement Requirements

Any department wishing to acquire an AI tool must submit a written request to the IT Director (or designated lead) describing: the clinical or operational purpose; data involved, including PHI; proposed risk tier and security level; BAA status; and for Tier A tools, clinical validation evidence.

Fill in the approval authority sign-off fields for Tier A tools. These require written approval from clinical, IT, compliance and final approving authority before procurement is finalized.

16.2 AI Use Register

The AI Use Register (Appendix B) is the central accountability mechanism. If a tool is not in the Register, it should not be in use. Fill in the role responsible for maintaining Appendix B — typically the IT Director. The IT Director updates Appendix B whenever a tool is added, retired or reclassified and notifies the CCO within five business days of any change.

16.3 Procurement Preferences

Select all evaluation criteria and procurement preferences that your hospital will consider when assessing AI tools and systems. These criteria are intended to support risk-informed procurement, vendor review and implementation decisions. While not all criteria may be required for every AI application, they should inform the hospital's evaluation process based on the level of clinical, operational, privacy, security and patient safety risk associated with the proposed use case.

Ask vendors to disclose known performance limitations in rural or low-volume settings as part of standard procurement due diligence, using the same disclosure process used for other bias and limitation reviews (see the Applied Model Card / AI Vendor Questionnaire referenced in *Figure 7*, page 33). This is especially relevant for hospitals in rural service areas, where population-specific validation is less likely to exist by default.

16.4 Decommissioning

Decommissioning decisions for Tier A tools require written authorization from the CMO and notification to the CCO. For Tier B and C tools, the IT Director may authorize decommissioning with notification to the CCO within five (5) business days. If an AI tool's performance degrades or poses a patient safety concern, the CMO may order immediate suspension pending review; the IT Director may order immediate suspension for cybersecurity or technical integrity concerns, with CMO notification within 24 hours.

All decommissioning actions must be reflected in the AI Use Register (Appendix B) within five (5) business days.

When an AI tool or system is retired, hospitals should complete all applicable decommissioning requirements prior to discontinuing use. At minimum:

- PHI, clinical data and related records should be reviewed for appropriate archival, transfer, retention or secure destruction in accordance with HIPAA, organizational records retention policies and applicable legal requirements
- The hospital's AI inventory or AI Use Register should be updated to reflect the tool's retired status
- Personnel, departments and affected users should be notified to discontinue use of the tool
- User access, integrations and system connections should be disabled or removed as appropriate

Hospitals should not overlook PHI and data management obligations during decommissioning, as improper handling of retained data, archived records or residual system access may create privacy, security, compliance and operational risks.

Section 17: Incident Response and Occurrence Reporting

Estimated Time 30-45 minutes

Section 17 integrates AI incident reporting with your hospital's existing occurrence reporting system and HIPAA breach notification process. This section requires the most hospital-specific fill-ins of any section in the *AI Use Policy Template*.

The incident response table in Section 17 has four rows — one for each event type. For each row, fill in: (1) the name of your occurrence reporting system (e.g., RL Solutions, Quantros, Datix); and (2) the specific names or contact information for the responsible parties listed.

Figure 5. **Organization-Specific Incident Response Contacts and Policy References**

Event Type	What to Fill In
Clinical AI event with potential patient harm (Tier A)	Name of your occurrence reporting system; charge nurse contact; CNO name and extension
Suspected HIPAA breach or unauthorized PHI disclosure	IT Director name and extension; Privacy Officer name and extension; Breach Notification Policy number
Cybersecurity incident	IT Director name and extension; Cybersecurity Incident Response Plan policy number
Policy violation or personnel misuse	CCO name and extension; Ethics Hotline number

Do not leave any contact field blank in Section 17. Staff facing an AI-related incident need to know exactly who to call and how. A blank field in an incident response table means the escalation path has a gap — and that gap will appear at the worst possible moment.

Section 18: Patient Rights and AI Disclosure

Estimated Time 15-20 minutes

Section 18 establishes the patient's right to information about AI use in their care. This is supplementary to your Patient Rights and Responsibilities Policy. Check the cross-reference policy number field so surveyors can navigate between the two documents.

Check all patient rights commitments that apply. At minimum: patients are informed when AI materially influences a diagnosis or treatment recommendation, and patients may request a plain-language explanation of how AI was used in their care.

The cross-reference to your Patient Rights and Responsibilities Policy (internal policy number) is important for Joint Commission survey readiness. Surveyors reviewing your AI policy should be able to find the corresponding patient rights document. **Fill in the policy number — do not leave it blank.**

Section 19: Policy Review and Effective Date

Estimated Time 15-20 minutes

The AI landscape changes quickly. A policy that accurately reflects your hospital's practices today may need updating within a year due to new tools, new regulations or significant changes in operations. The review cycle in Section 19 establishes a regular schedule for keeping the policy current.

19.1 Review Cycle

Annual review is the standard recommendation. Check annually unless your hospital has a specific reason to review more or less frequently. Whatever you choose, the policy should also be reviewed whenever there is a significant change in: AI tool inventory, applicable law or regulation, accreditation standards or a significant AI-related incident.

19.2 Policy Approval Signatures

Section 19.2 contains the Policy Approval Signatures page. Every signature field should be completed before the policy takes effect. Approval roles should reflect the governance structure established by the organization (see Appendix A). In many hospitals, this includes the CEO or President, a clinical leader (such as the CMO) when clinical AI tools are covered, a compliance or privacy leader, and an IT security leader responsible for reviewing cybersecurity and technical safeguards. Organizations may adapt approval roles to align with their size, staffing structure, and governance model. Board ratification is recommended and may be required by some accrediting organizations.

For CAHs and smaller community hospitals, a single individual may sign multiple approval lines if they hold multiple governance responsibilities. Any dual-role assignments should be documented clearly in Appendix A.

19.3 Staff Acknowledgment Form and 19.4 Record of Revisions

Section 19.3 is the Staff Acknowledgment Form. Retain completed acknowledgments in personnel files or your LMS. Section 19.4 is the Record of Revisions — update it every time the policy changes and treat it as your audit log. The effective date in the policy header is when all obligations take effect: confirm staff are trained, vendors are informed of the transition period, and all governance roles are documented in Appendix A before the policy goes live.

Part 2 — Add-On Modules

A: Ambient Documentation AI

Estimated Time 20-30 minutes

Include this section if your hospital uses any ambient documentation AI (e.g., Nuance DAX, Abridge, Suki) or any AI that generates content entered into the EHR medical record. This section is required if any Tier A AI tool produces clinical documentation.

A.1 Clinician Attestation and Medical Record Finalization

No additional fill-in is required in Section A.1 — the policy language is complete as written. This provision reflects the fundamental principle of medical record law: the signing clinician is fully accountable for the content of the finalized record, regardless of how it was generated. This should be highlighted in training for all clinical staff who use ambient documentation tools.

A.2 Required Patient Disclosure for Ambient Documentation Tools

List all approved ambient documentation tools in the fill-in table. Then go to Appendix C to complete or add the approved disclosure script for each tool. The CMO must approve each script before it is used with patients. Clinical staff may not improvise their own disclosure language.

Document patient declination in the EHR using consistent, searchable language. A suggested format: “Patient declined AI-assisted documentation. Visit documented manually per patient request.” This notation creates an auditable record and demonstrates compliance with patient rights obligations. obligations take effect: confirm staff are trained, vendors are informed of the transition period, and all governance roles are documented in Appendix A before the policy goes live.

> Ambient documentation tools trained primarily on general American English speech patterns may perform less reliably with regional dialects, local speech patterns or lower-bandwidth audio common in some rural clinical settings. Ask the vendor whether accuracy testing included populations and recording conditions comparable to your Hospital’s patient population.

A.3 Governance and Oversight

Cloud-based ambient documentation tools generally require a stable internet connection to function. Hospitals in areas with inconsistent broadband should confirm the tool’s behavior during a dropped or interrupted connection before deployment and ensure staff have a clear, practiced fallback to manual documentation, so patient care isn’t disrupted.

B: Diagnostic and Radiology AI

Estimated Time 25-35 minutes

Include this section if your hospital uses any AI tool that reads, flags, triages or prioritizes imaging, pathology or laboratory results. Common examples include radiology AI that worklists CT scans for suspected stroke or pulmonary embolism, AI-assisted pathology slide reading, and EHR-integrated lab value flagging tools.

B.1 Clinical Validation Requirements

Before completing this section, confirm whether your hospital has already conducted a formal validation review for each diagnostic or radiology AI tool in use. If a tool was deployed without a documented

validation, this section creates the opportunity to require one retroactively.

Select all check boxes that reflect your hospital's validation requirements. At minimum, most hospitals should check the first four items:

- Vendor evidence of training data and performance characteristics is the baseline for any diagnostic AI procurement decision.
- Known performance limitations must be documented so clinical staff understand when the tool may be less reliable.
- Local validation or pilot testing is strongly recommended, particularly for tools used in populations that differ from the vendor's training data.
- Bias assessment across race, ethnicity, gender and age is required for tools used in clinical decision-making under CMS and Joint Commission guidance.

For hospitals serving rural or critical access populations, ask vendors directly whether their training and validation data included rural or low-volume settings. Many diagnostic and predictive AI tools are trained predominantly on data from larger, urban health systems. A hospital cannot assume equivalent performance without evidence and should document this gap the same way it documents any other known limitation.

Assign the validation review to the CMO or a CMO-designated clinical lead with appropriate expertise in the tool's diagnostic domain. For radiology AI, this is typically the Chief of Radiology. Fill in the "Validation review assigned to" field with the role title, not a name, so the policy does not require revision when personnel change.

B.2 Human Oversight Requirements

Select all check boxes that apply. The first three reflect standard-of-care requirements for diagnostic AI in U.S. clinical settings and should be checked by virtually all hospitals:

- All AI-generated diagnostic outputs must be reviewed by a qualified clinician before influencing patient care. This is non-negotiable for Tier A tools and reflects both patient safety and legal accountability standards.
- The radiologist or pathologist of record retains full professional and legal accountability for finalized interpretations. This should be reinforced in medical staff training.
- AI outputs may not be used to delay or override clinical judgment. If a clinician's assessment differs from the AI output, the clinician's judgment governs.

B.3 Approved Diagnostic AI Tools

Do not list tools in the body of this section. Go to Appendix E — Approved Tool Lists and complete the Module B table there. For each approved tool, record the tool name and vendor, diagnostic function, BAA status and date, CMO validation date, and oversight owner by role title. Update Appendix E when tools are added, retired or replaced without revising the core policy. The tools should also be registered in the AI Use Register (Appendix B).

B.4 Medical Staff Oversight

This section mirrors the Medical Staff credentialing language in Module A. Select the check boxes that match your Medical Staff Bylaws and governance structure. If your hospital's Medical Executive Committee (MEC) has not yet established a process for reviewing diagnostic AI, use this policy section to create that expectation. Reference the relevant Bylaw section or policy number in the fill-in field.

If any diagnostic AI tool at your hospital produces interpretations without radiologist or pathologist review — even as a "preliminary" read — the check box requiring a specific clinical privilege is strongly recommended. Autonomous diagnostic AI sits at the highest risk tier and warrants explicit Medical Staff credentialing oversight.

C: Predictive Alerts and Clinical Decision Support

Estimated Time 20-30 minutes

Include this section if your hospital's EHR or any connected system generates AI-driven clinical alerts, risk scores or deterioration flags that appear in nursing or physician workflows. This includes sepsis prediction models, early warning systems, readmission risk flags and similar tools.

C.1 Alert Fatigue and Clinical Safety

Alert fatigue is the desensitization of clinical staff to high volumes of alerts, including AI-generated ones, and is a documented patient safety risk. This section establishes joint CMO and CNO accountability for monitoring alert impact.

Select the check boxes that match your hospital's current or intended oversight practices. The first two are strongly recommended for any hospital using active AI alert tools:

- Semi-annual review of alert volume and clinical response rates provides early detection of alert fatigue before it becomes a patient safety event.
- CMO-led review triggered by alert fatigue findings ensures escalation is structured and documented, not ad hoc.

Fill in the alert threshold review role by title — typically the CMO, Chief Medical Information Officer (CMIO) or a designated quality officer. This person is responsible for reviewing AI alert sensitivity settings before any changes are made. Do not assign this to an IT or vendor contact; it must be a clinical governance role.

> Predictive alert tools such as sepsis or deterioration models are often calibrated using response-time assumptions drawn from larger, better-staffed systems. A rural or critical access hospital with fewer staff on a given shift may not be able to respond to alerts at the pace the tool assumes. Alert fatigue and safety review should account for actual available staffing, not only alert volume.

C.2 Approved Predictive Alert Tools

List all approved predictive alert tools in Appendix E — Module C. For each tool, include the tool name and vendor or EHR module name, alert type, BAA status, alert response owner (the clinical role responsible for acting on the alert), and date of last threshold review.

C.3 Clinical Decision Support (CDS) Human Override Requirements

All three check boxes in this section should be selected by most hospitals. They reflect core principles of safe clinical decision support deployment:

- Clinicians must be able to override any AI alert without penalty. Policies or workflows that penalize override — explicitly or implicitly — undermine clinical judgment and create liability.
- Override pattern review by the CMO turns individual clinical decisions into a population-level safety signal. If a large proportion of clinicians are overriding a specific alert, the threshold may need adjustment.
- AI risk scores are one input, not a clinical determination. This language is critical for staff training and should be reinforced in any AI tool onboarding.

C.4 Incident Reporting for Alert-Related Events

Fill in the name of your hospital's occurrence reporting system (e.g., RL Solutions, Quantros, Origami, Meditech Safety) and the CNO notification contact for Tier A alert events.

Any clinical situation where an AI-generated alert may have contributed to patient harm — or where a clinician's failure to act on an alert contributed to harm — should be filed immediately. When in doubt, file. The occurrence report creates a record that enables the Patient Safety and Quality Committee to assess whether the alert design, threshold or workflow requires modification.

D: Revenue Cycle AI

Estimated Time 20-30 minutes

Include this section if your hospital uses AI for any function that affects Medicare or Medicaid billing — including coding assistance, code review, claims editing, prior authorization support or denial management. This module is required if any AI tool generates or modifies codes submitted to federal payers.

False Claims Act risk: AI-assisted coding errors that result in overbilling to Medicare or Medicaid can constitute false claims under 31 U.S.C. § 3729. The False Claims Act imposes civil penalties and treble damages. OIG has identified AI-assisted upcoding as an active area of concern. Human review of all AI-generated codes before submission is not optional.

D.1 Revenue Cycle AI Requirements

Select all check boxes that apply. The first check box — requiring a qualified coder or clinician to review and approve all AI-generated codes before submission — is mandatory for any hospital subject to Medicare or Medicaid billing rules. It should not be unchecked.

For the remaining check boxes:

- Prior authorization tools: If your hospital uses AI to generate or populate prior authorization requests, check this box. The final authorization decision must comply with payer contracts and CMS requirements regardless of AI involvement.
- Compliance Program audit schedule: Revenue cycle AI should be treated as a billing risk area and included in your annual compliance audit plan. If your hospital has a Corporate Compliance Officer (CCO), coordinate with them to add AI tools to the audit schedule.
- Vendor attestation: Require in writing that revenue cycle AI vendors attest their tools comply with applicable CMS billing rules and OIG guidance. This should be a standard contract term, not a verbal assurance.
- Upcoding/undercoding review: If the AI tool flags unusual billing patterns, those findings must be reviewed by the CCO and reported through the Compliance Program — not suppressed or ignored.

Assign oversight to the CCO or Revenue Cycle Director by title. This role is responsible for ensuring the requirements in this section are operationalized and for reporting AI-related billing findings to leadership. Do not assign to a vendor contact or IT role.

> Rural patients often have fewer alternate providers within reasonable travel distance, so a denied or delayed prior authorization can carry a heavier practical burden than it would for a patient with more options nearby. Fold this consideration into the CCO's existing review of AI-identified upcoding, undercoding or unusual billing patterns. A separate process is not necessarily needed.

D.2 Approved Revenue Cycle AI Tools

Complete the Module D tool list in Appendix E. For each tool, document the tool name and vendor, function, BAA status and date, coder/clinician review process and oversight owner by role title.

D.3 Board Oversight

This section establishes Board of Trustees accountability for AI governance — the appropriate level of oversight for tools that carry False Claims Act and HIPAA risk. Select the check boxes that match your Board's current governance practices and risk tolerance. Most hospitals should check the first three:

- Annual AI governance report to the Board from the CCO provides the Board with the visibility needed to fulfill its fiduciary responsibility for compliance risk.
- Notification of significant patient safety events or HIPAA breaches involving AI ensures the Board is not the last to know when something goes wrong.
- Board ratification of this policy and major revisions signals that AI governance is a Board-level priority, not solely a management function.

Specify the role presenting the annual AI report (typically the CCO or CEO) and the frequency and Board meeting at which it will be presented (e.g., "Annual — January Board meeting"). This creates a standing agenda item that keeps AI governance visible at the governance level.

E: Generative AI and Productivity Tools

Estimated Time 20-30 minutes

Include this section if any staff member uses ChatGPT, Microsoft Copilot, Google Gemini, or similar generative AI tools for work purposes — including on personal devices. In practice, most hospitals should include this module. Consumer generative AI use by clinical and administrative staff is widespread, and the absence of explicit policy coverage creates significant HIPAA and compliance exposure.

E.1 Authorized Use of Generative AI

The check boxes in this section link authorized generative AI use to the hospital's risk tier framework from Section 10. Before completing this section, confirm that your hospital's AI Use Register (Appendix B) is current and includes all known generative AI tools in use.

Select all check boxes that apply:

- Tools must appear in the AI Use Register with an assigned tier and security level before staff may use them for work purposes. Unapproved tools default to prohibited.
- Tier B tools (e.g., enterprise-licensed Microsoft Copilot for M365) require a Data Processing Agreement or BAA if PHI may be involved. Verify BAA status before checking this box.
- Tier C tools (e.g., free-tier ChatGPT) may only be used for internal, non-sensitive, non-PHI tasks. Outputs are informal drafts and must be treated as such.
- Staff must complete required AI training before first use of any authorized generative AI tool. Confirm your training program is in place before activating this check box.

A staff member's personal paid subscription — such as a personal ChatGPT Plus or Claude Pro account — is classified as Level 3 even if the employee pays for it. These tools carry no Hospital data protections and may only be used for Tier C tasks with no PHI or PII. If your hospital prefers to prohibit personal AI tools entirely for work-related tasks, use the optional prohibitory language provided in the *AI Use Policy Template*.

E.2 Prohibited Generative AI Use

Select all check boxes that apply. The first four reflect prohibitions required to maintain HIPAA compliance and protect patient and staff privacy:

- PHI into non-BAA tools: Entering any patient information — including names, dates of service, diagnoses or any combination of identifiers — into a publicly accessible AI tool is a HIPAA violation. This prohibition must be reinforced in staff training and should be one of the first topics covered in AI onboarding.
- Misrepresentative content: AI-generated content that falsely attributes statements to hospital leadership, clinicians or the hospital itself creates legal and reputational risk. All AI-generated external communications must be reviewed before publication.
- Deepfakes and synthetic media: Explicitly prohibit the generation of any content that impersonates a clinician, patient or hospital official. This includes AI-generated voices, images or video.
- Patient-facing clinical substitution: AI-generated responses may not substitute for qualified clinical advice in any patient-facing context — including patient portal messages, after-hours messaging or discharge instructions.

E.3 Approved Generative AI Tools

Complete the Module E tool list in Appendix E. For each approved tool, document the tool name and vendor, authorized use cases (be specific — e.g., "draft internal communications, meeting summaries, non-clinical administrative documents"), risk tier and security level, BAA or Data Processing Agreement (DPA) status, and oversight owner by role title. Review and update this list at least annually, or when new generative AI tools are adopted or discontinued.

E.4 AI-Generated Content Review Requirements

Select all check boxes that apply. These requirements govern the quality and accountability of AI-generated outputs used in hospital operations:

- Official communications and patient education: Any AI-generated content used in external-

facing materials — including patient education handouts, website content, press releases or community communications — must be reviewed and approved by a qualified staff member before publication.

- Clinical documentation attestation: Staff may not submit AI-generated clinical documentation, reports or analyses without personal review and attestation. This mirrors the attestation requirement in Module A and applies to any AI-generated clinical content, not only ambient documentation.
- Attribution and disclosure: AI-generated content must not be represented as the original work of a specific individual without disclosure. This applies to publications, grant applications and any content where authorship is material.

Module E affects the broadest range of hospital staff. When rolling out this policy, prioritize training that covers the PHI prohibition (E.2), the tier framework for authorized tools (E.1) and the content review requirements (E.4). A one-page quick reference card summarizing what staff can and cannot do with AI tools is a practical complement to the formal policy.

Generative AI increasingly powers patient-facing communication such as appointment reminders, portal messages and chatbots. For hospitals serving rural populations, confirm these tools don't become the only channel. Patients without reliable broadband or smartphone access need an equivalent non-digital option.

App. A: Appendix A — Governance Roles and Contacts

Estimated Time 20-30 minutes

Complete the Appendix A table by assigning a specific person by name, title and contact to each governance role. Appendix A may be updated when personnel change without requiring a formal policy revision — but every update must be communicated to the policy owner within five (5) business days.

The last column — Contact — should include at minimum a phone extension or email address. Staff need to be able to reach the right person quickly during an incident. A list of names without contact information does not serve that purpose.

The rural/community AI liaison role is optional and is intended for hospitals that already participate in a regional or community advisory structure (e.g., a hospital community benefit committee). Where such a structure exists, naming a liaison here creates a clear channel for routing high-impact AI tool decisions for community input, consistent with Section 7's governance discussion. Hospitals without an existing structure should leave this row blank rather than create new infrastructure solely to fill it.

App. B: Appendix B — AI Use Register

Estimated Time 30-45 minutes

The AI Use Register is the central accountability mechanism of the policy. If a tool is not in the Register, it should not be in use. Appendix B should be populated before the policy's effective date — staff cannot comply with a policy that tells them to check the Register if the Register is empty.

Getting Started with the AI Use Register

Before completing Appendix B, conduct a brief AI inventory — ask department heads and IT what tools are currently in use that may involve AI. This is a suggested starting question for each department:

"What tools do you use that automatically generate recommendations, flags, scores, summaries or drafts?" Common surprises in hospital AI inventories include:

- EHR-embedded predictive models (e.g., Epic's deterioration index, readmission risk flags)
- Radiology worklist prioritization tools
- Ambient documentation pilots that may have started as informal trials
- Revenue cycle coding tools with AI-assisted features
- Staff scheduling optimization software with predictive components

For each tool identified, fill in: Tool/Vendor, Purpose, Risk Tier, Security Level, BAA/Agreement status and reference number, Oversight Owner, Date Approved and Date Last Reviewed. No tool should appear in the Register without a completed Tier assignment and BAA status.

Also note whether the tool has been validated for rural or low-volume populations comparable to this Hospital's service area. Where no such validation exists, record that explicitly rather than leaving the field blank: an undocumented gap is harder to catch during audit or review than a documented one.

The person who confirms the "rural / low volume validation status" field in Appendix B should be the same role assigned to the rural/low-volume validation disclosure control in Appendix F (see item 19). Splitting ownership between the two appendices creates a gap where the Register looks complete but no one actually verified the underlying disclosure.

AI Use Register Maintenance Tip

Designate the IT Director as the Register owner and build in a quarterly calendar reminder to review and update it. Tools go out of use, vendors change their terms of service, and new tools get adopted informally. A stale Register is almost as dangerous as no Register — it creates a false sense of control. Quarterly review keeps it accurate.

Rural Validation Maintenance Tip

When adding a new tool to the Register, capture the vendor's answer on rural/low-volume validation at the same time as BAA status and risk tier. Don't leave it for a later review. If the vendor cannot answer, record "Not validated, documented limitation" rather than leaving the field blank. Revisit this field on the same quarterly cadence recommended for the rest of the Register.

App. C: Appendix C — Approved Patient Disclosure Scripts

Estimated Time 20-30 minutes

Appendix C contains the CMO-approved verbal disclosure scripts for each ambient documentation or patient-facing AI tool. Completing this appendix is a prerequisite to clinical use of any ambient documentation tool — no script means no authorized use.

Where a hospital serves patients who travel significant distances or face limited follow-up access, disclosure scripts should also make clear that declining an AI tool does not delay, reduce or otherwise change the care they receive that visit. This reduces the risk that patients with fewer opportunities to return will decline an AI-assisted tool out of concern it is required to receive care.

Completing a Disclosure Script Entry

For each tool, complete: the CMO approval date, the CMO's name as approving authority, the applicable tool name, and the verbatim approved script text. The sample script language in the *AI Use Policy Template* is a starting point — adapt it to

your hospital's communication style and each tool's specific function.

Key elements to preserve in any disclosure script:

- Name of the AI tool or describe what it does
- Confirmation that all AI-generated notes will be reviewed and approved by the clinician before entering the chart
- Offer of the right to decline

For hospitals in rural service areas, consider whether the script should explicitly address travel distance or limited opportunities for follow-up. Patients who traveled 45 minutes for an appointment may feel greater pressure to accept a tool they are uncertain about than patients who can easily return for another visit. The script should make the decline option feel free of consequence.

Patient declination must be respected immediately and documented in the EHR. Suggested EHR notation: '[Tool name] declined by patient. Visit documented manually per patient request.' No adverse inference may be drawn from a patient's decision to decline AI-assisted documentation.

App. D: Appendix D — Regulatory References and Internal Policy Cross-References

Estimated Time 20-30 minutes

Appendix D maps each regulatory citation to your hospital's corresponding internal policy number. This is the cross-reference document surveyors and auditors use to navigate your policy ecosystem. A well-completed Appendix D signals organizational maturity and survey readiness.

For each row in the table, enter your hospital's internal policy number in the right-hand column. If a listed item does not yet have a corresponding internal policy, note "In development — target date: ___" rather than leaving it blank. Gaps are acceptable; undocumented gaps are not.

Rural proofing resources are not regulatory requirements and don't require a hospital policy number. These resources allow for cross-reference so reviewers can see that rural-specific guidance was consulted during policy development.

App. E: Appendix E — Approved Tool Lists by Module

Estimated Time 15-20 minutes

Appendix E contains the module-level tool inventories. Complete only the sections for modules your hospital has selected (E-A through E-E). All tools listed in Appendix E must also appear in the AI Use Register (Appendix B) — they serve different purposes: Appendix B is the master inventory; Appendix E captures the module-specific operational detail needed for audit and accreditation review, including clinical purpose, BAA references, CMO validation dates, safeguards and human review processes.

Carry the rural/low-volume validation status (e.g., "Validated for comparable population," "Not validated, documented limitation," "Not applicable to this tool's function") recorded in Appendix B through to each module-level entry in Appendix E. This keeps the detail visible at the point where clinical, compliance or accreditation reviewers are most likely to be looking, module by module, rather than requiring them to cross-reference back to the master register for this specific detail.

Update Appendix E when tools are added, retired or changed without requiring a formal policy revision. Communicate all updates to the Policy Owner within five (5) business days. Section E-A.1 includes a Tier A documentation safeguards table — document the clinical purpose, PHI status, safeguards and final decision authority for each ambient documentation tool there.

App. F: Appendix F — Control Role Assignments

Estimated Time 20-30 minutes

Appendix F assigns the responsible role for each required control under the risk tier framework from Section 10. It is organized by tier. For Tier A tools, assign roles for: data security and BAA review, clinical validation before deployment, human-in-the-loop licensed clinician review, incident notification and rural/low-volume disclosure review within 24 hours. This control operationalizes the rural-proofing review described in Module B.1 and Section 16.3 by assigning clear ownership: someone must confirm, before a Tier A tool goes live, that the vendor's rural/low-volume validation disclosure (or the documented absence of one) is on file, not just that it was asked for.

Like Appendix A, Appendix F may be updated when roles or personnel change without a formal policy revision. Communicate all updates to the Policy Owner within five (5) business days. In CAHs and smaller hospitals where one person holds multiple roles, document those assignments clearly and identify a backup for critical Tier A control functions.

Common Pitfalls to Avoid

Implementing an AI governance program involves more than adopting a policy. Many challenges arise during implementation when governance processes, accountability structures, and operational procedures are not fully established. Figure 6 highlights common pitfalls hospitals may encounter when implementing this policy and provides practical strategies to help organizations avoid them. Reviewing these considerations before policy adoption can help support smoother implementation and ongoing compliance.

Figure 6. **Common AI Governance Pitfalls and How to Avoid Them**

Pitfall	What to Do Instead
Leaving Appendix B (AI Use Register) empty when the policy takes effect	Conduct a brief AI inventory before the effective date. Staff cannot comply with a policy that references a Register that does not yet exist.
Treating AI policy as an IT project	AI governance is a clinical, compliance and organizational capability — not just a technology initiative. The CMO and CCO must be as engaged as the IT Director.
Forgetting AI embedded in the EHR	Epic, Cerner and other EHR systems have active AI modules that most hospitals have never formally classified. Ask your EHR administrator what predictive or AI-assisted features are currently enabled.
Clinical staff using consumer AI tools for patient documentation	Train staff explicitly: No PHI may ever enter a publicly accessible or non-BAA-covered AI tool. Treat any occurrence as a potential HIPAA breach.

Figure 6 (continued). **Common AI Governance Pitfalls and How to Avoid Them**

Pitfall	What to Do Instead
Skipping Medical Staff consultation for Modules A and B	Module A and Module B directly implicate Medical Staff Bylaws and credentialing. Bring the Medical Executive Committee (MEC) into the drafting process before the policy is finalized, not after.
Leaving contact fields blank in Section 17 (Incident Response)	Blank contact fields in incident response tables mean staff cannot escalate when something goes wrong. Every contact field must be filled in before the policy takes effect.
Writing a policy nobody reads or follows	Remove sections that do not apply to your hospital. A shorter, accurate policy that is followed is more valuable than a comprehensive one that is not. Tailor before you publish.
Forgetting to notify vendors of the transition period	Section 14.3 establishes a transition period for existing vendors. Notify vendors in writing within 30 days of policy adoption. Their clock starts when they are informed, not when the policy is signed.
Letting governance lapse after initial rollout	Schedule the next annual policy review before the current one is finished. AI tools and regulations evolve quickly. Set a calendar reminder for the CCO on the day the policy is approved.

Resources

Hospitals may benefit from a variety of resources when developing and implementing AI governance approaches. Figure 7 highlights selected frameworks, guidance documents, toolkits, and reference materials that may support risk management, compliance, vendor review, clinical oversight, and governance activities. This is not an exhaustive list. Organizations should consider additional resources as AI technologies, regulations, and industry practices continue to evolve. The list is organized alphabetically by the title of the resource.

Figure 7. **Selected Resources for Developing and Implementing Hospital AI Governance**

Title	Organization	Description	Link
AI in Rural Health Care: A framework for responsible adoption, readiness, and access	Journal of Rural Health	Rural AI challenges and opportunities; proposes a five-pillar Rural AI Readiness Framework	https://onlinelibrary.wiley.com/doi/10.1111/jrh.70144
AI in Rural Health Care: A policy roadmap for ensuring ethical and equitable usage and closing gaps in access	National Rural Health Association	Policy roadmap on rural AI challenges/ opportunities with federal, state and local recommendations across a five-pillar readiness framework	https://www.ruralhealth.us/nationalruralhealth/media/documents/advocacy/2026/nrha-artificial-intellegence-policy-paper.pdf
AI Risk Management Framework (AI RMF 1.0)	National Institute of Standards and Technology (NIST)	The federal standard for AI risk management, including governance, measurement and management of AI risk across an organization. Useful for grounding your risk tier framework.	https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
Applied Model Card (AI Vendor Questionnaire)	Coalition for Health AI (CHAI)	A standardized disclosure template vendors complete covering training data, validation, intended use and limitations. Use during vendor due diligence and Tier A tool approval review.	https://www.chai.org/workgroup/applied-model
Artificial Intelligence in Software as a Medical Device	U.S. Food and Drug Administration (FDA)	FDA's hub for AI/ML-enabled medical device guidance, including the SaMD Action Plan, Good Machine Learning Practice principles and predetermined change control plans. Relevant when evaluating whether a clinical AI tool is FDA-regulated.	https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-software-medical-device
AI Policy Template and Guidance	Kansas Health Institute (KHI), Health Resources in Action (HRiA) and Wichita State University Community Engagement Institute (WSU CEI)	Provides a framework for developing AI policies within public health organizations, including key considerations, implementation guidance and sample policy provisions.	https://www.khi.org/articles/developing-artificial-intelligence-ai-policies-for-public-health-organizations-a-template-and-guidance/
Compliance Guidance	U.S. Department of Health and Human Services (HHS) Office of Inspector General (OIG)	OIG guidance on AI in coding and billing.	https://oig.hhs.gov/compliance/compliance-guidance/

Title	Organization	Description	Link
Gaps in AI Research for Rural Health in the U.S.: A scoping review	JAMIA	Scoping review of 26 studies; finds AI research rarely reaches deployment and no generative AI use in rural settings	https://academic.oup.com/jamia/article/33/2/509/8341616
Governance for Augmented Intelligence Toolkit	American Medical Association (AMA), with Manatt Health	An eight-step playbook for standing up AI governance, including a downloadable model policy, vendor-vetting worksheets and sample forms.	https://edhub.ama-assn.org/steps-forward/module/2833560
Health System AI Governance (white paper)	Duke-Margolis Institute for Health Policy	A multi-health-system research project comparing how organizations of varying sizes have structured AI governance in practice.	https://healthpolicy.duke.edu/sites/default/files/2024-10/AI%20Governance%20in%20Health%20Systems.pdf
Healthcare Exceptionalism in Patient Perspectives on AI	Health Affairs Scholar	National survey finds patients are less comfortable with healthcare AI than with familiar everyday AI (e.g., GPS, fraud alerts, etc.)	https://academic.oup.com/healthaffairsscholar/article/4/6/qxag128/8690430
Responsible Use of AI in Healthcare (RUAIH)	Joint Commission and Coalition for Health AI (CHAI)	A national accreditation-adjacent framework covering seven core areas of hospital AI governance (policy/governance structure, risk and bias assessment, monitoring, training), underpinning a voluntary Joint Commission certification.	https://www.jointcommission.org/en-us/certification/responsible-use-of-ai-in-healthcare
Rural Proofing Artificial Intelligence	Federation of American Scientists	Framework and five-step action plan for ensuring AI tools/policies reflect rural realities and avoid unintended harm	https://fas.org/publication/making-rural-communities-visible-ai/
Sample Business Associate Agreement Provisions	HHS	Model BAA provisions from HHS that can be adapted when procuring AI tools that handle PHI. Use as a baseline for vendor BAA negotiations.	https://www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions/index.html
Trustee Insights: AI Governance for Hospital Boards	American Hospital Association (AHA)	Resources for preparing Board-level AI governance reports.	https://trustees.aha.org/artificial-intelligence-and-trustee
Unintended Consequences of Using Ambient AI Scribes for Billing	JAMA Health Forum	Warns that using ambient scribes to drive higher-intensity billing risks increased spending, clinician strain and eroded patient trust	https://jamanetwork.com/journals/jama-health-forum/fullarticle/2843723


KANSAS HEALTH INSTITUTE
Informing Policy. Improving Health.


HEALTH
WORKS
KANSAS HOSPITAL
ASSOCIATION