COMMUNITY HEALTH CARE ASSOCIATION of New York State

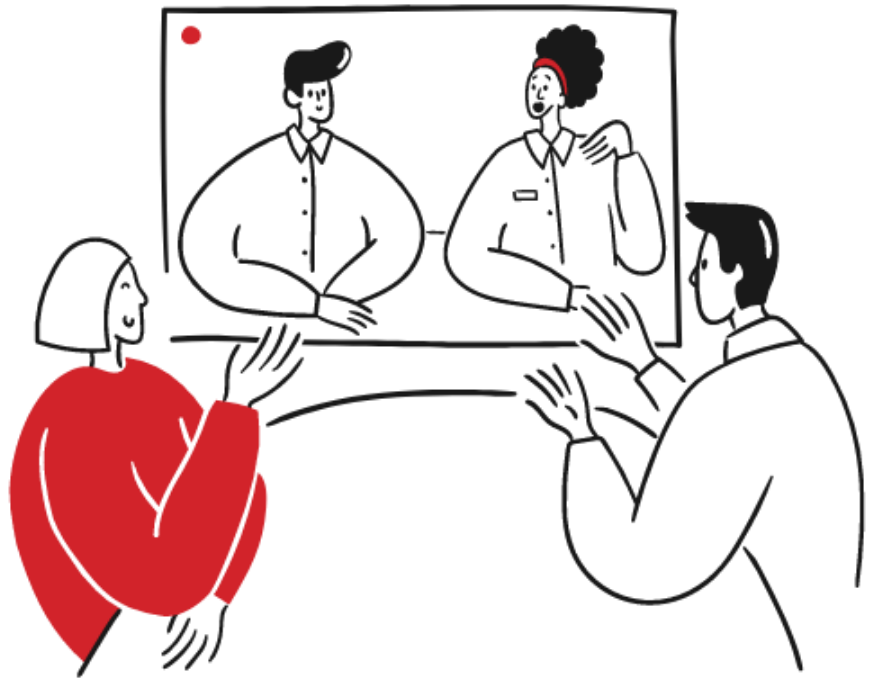
A CHCANYS NYS-HCCN funded webinar with CHAI:

Building a Foundation for Responsible Super Intelligence: An Introduction to CHAI's Policy Playbook

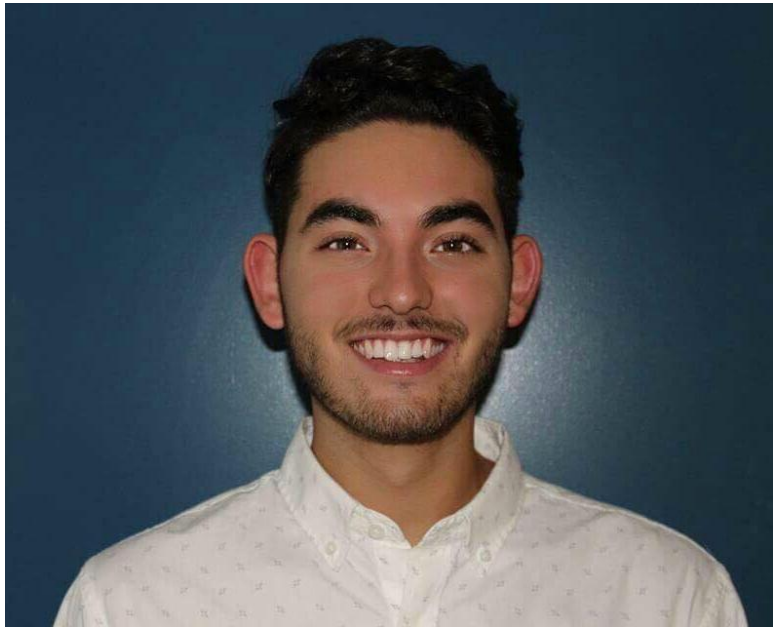
November 10, 2026

Zoom Guidelines

- Please share your questions in the chat. CHCANYS staff will raise your questions to our speakers and follow up as needed if there are unanswered questions.
- The session is being recorded and materials will be shared after the session.



Meet the Speakers



Anthony DiDonato

Program Manager, CHAI



Merage Ghane, PhD

Director of Responsible AI, CHAI

CHAI

Domain 1: SI Policy Playbook

Where SI governance starts



Coalition for Health AI | November 2026

Educational Webinar • 60 min

Note:

HRSA has directed HRSA-funded entities (like CHCANYS) to replace the term “artificial intelligence (AI)” with “super intelligence (SI)” in all federally funded communications and work, consistent with the President’s Executive Order. At this time, SI is synonymous with AI. SI will be primarily used in the presentation and AI will be used when referencing CHAI materials/playbooks or other proprietary named documents/resources.

Agenda

01

~5 min

Where are you today?

A quick poll on how familiar you are with your organization's SI policy — and how you use SI now

02

~5 min

AI governance, and where policy fits

The CHAI framework's four domains, and why Domain 1 is the entry point

03

~20 min

The three SI policy controls

Each control in plain language, how a community health center could implement it, with notes, examples and references

04

~30 min

Build it together

An interactive walk through the Quick-Start handout: what to consider at every checklist item, your red lines, and your first 30 days

INTERACTIVE

LIVE POLL

How familiar are you with your organization's SI policy?

Q1 — Which best describes you?

- **A** We don't have one — or I don't know if we do
- **B** We have one, but I haven't read it
- **C** I've read it and know roughly what it covers
- **D** I helped write it, or I maintain it

Q2 — How do you use SI at work today?

- **1** I don't use SI at work
- **2** I use consumer tools (e.g., ChatGPT, Copilot)
- **3** I use SI features built into our EHR or other systems
- **4** I'm involved in selecting, building, or governing SI tools

STOP 02 • ~7 MINUTES

SI governance — and where policy fits

The CHAI SI Governance Framework's four domains, and why Domain 1 is the entry point

Four domains: policy is the foundation

Domain 04 Organizational Processes

How: lifecycle management (4.1), risk & impact assessments (4.2), data management (4.3), third-party management (4.4), education & training (4.5)

Domain 03 Organizational Resources

With what: the people, budget, and tooling to run governance

Domain 02 Organizational Structures

Who: governance committee, roles, reporting lines

Domain 01 AI Policy

How do we govern AI, what is in scope, and who is accountable?

TODAY

Note: how to read the playbooks

Baseline operational controls were consensus-defined by healthcare delivery organizations as 'must-haves'. The implementation guidance is suggestions based on known practices — implementation will differ by goals, resources, and maturity.

Example: who built them

More than 150 organizations across the healthcare ecosystem contributed — health systems, clinicians, technologists, payers, patient advocates, professional societies.

Reference

CHAI AI Governance Playbooks: Introduction & Background; governance.chai.org

Without a policy, SI is being used, just ungoverned

Four doors SI walks through (with or without a policy)

Embedded in systems you already own

AI features switched on inside the EHR or patient portal with a vendor update

Vendor-supplied tools

Ambient documentation, call-center chatbots, coding and claims tools

Consumer and copilot tools

Staff using ChatGPT, Gemini or Copilot to draft, summarize, or answer questions

Built in-house

A team builds a small prediction or triage model with local data

Note: shadow SI

Any SI system, feature, or capability that is developed, procured, embedded, or used by staff without going through the organization's established SI intake, risk assessment, and approval processes.

For example:

- A department uses an unapproved recording/note taking app on their phone that uses generative SI to summarize patient notes
- A vendor turns on an SI feature inside the EHR without org determining if it needs separate review/training
- Staff upload documents or data into third-party SI solutions that hasn't been approved

The policy is what creates the 'established process' — it turns shadow SI into governed SI.

Recommended reading (playbook): 'Health Systems Only Govern the Tip of the Iceberg'

STOP 03 • ~17 MINUTES

The three SI policy controls

What each one asks of you, and how a community health center could put it into practice

Three baseline operational controls: the 'must-haves'

Control 1.1

The Foundation

"Develop a formal SI Policy document, approved by organizational leadership, and make accessible to all relevant staff."

In plain terms: Write it down, get it approved at the top, and make sure staff can find it.

Control 1.2

Key Concepts & Scope

"Define key concepts in the SI Policy (e.g., definition of SI, scope of SI solution use, SI approval process)."

In plain terms: Say what you mean: what counts as SI, what's in scope, how something gets approved.

Control 1.3

Changes & Maintenance

"Establish an appropriate individual or individuals with designated oversight for policy changes."

In plain terms: Name the person who owns keeping it current.

Reference

CHAI Domain 1: AI Policy Playbook — 'Baseline Operational Controls for Domain 1: AI Policy' (governance.chai.org/domain-1-ai-policy). Controls were consensus-defined by healthcare delivery organizations.

Write it down, get it approved, make it findable

THE CONTROL SAYS *"Develop a formal SI Policy document, approved by organizational leadership, and make accessible to all relevant staff."*
In plain terms: Write it down, get it approved at the top, and make sure staff can find it.

In practice — how a community health center could implement it

1. Secure executive commitment

One-on-ones with the CEO / CMO / CNO / CIO for written commitment and an initial budget or repurposed resources. Build the business case on risk mitigation, workforce protection and patient safety — not technology.

2. Obtain formal approval

Cross-functional review (compliance, IT, clinical, operations — often your existing compliance or quality committee) so the policy doesn't conflict with existing policies or state/federal rules. Approve at the Board or Executive Committee and log it.

3. Make it accessible and inform staff

Publish on the intranet and website; add key sections to role-specific training; name a champion to take it to department meetings; consider a staff attestation and a usage request form.

Note: securing executive buy-in

Executive investment can be hard to get — resource constraints, low SI literacy or trust, different motivators. Personalize the pitch to the executive's role and lead with organizational risk, not a pitch for new technology.

Example: five arguments that have worked

Lead with risk (patient safety, workforce, regulatory/liability, reputational, financial, bias) • real failure stories and near-misses • governance as faster, safer adoption • protecting clinicians and staff • ROI and value-on-investment tied to risk and sustainability.

Reference

Domain 1 Playbook — Control 1.1 implementation guidance, Tasks 1–3; notes on executive buy-in and shadow SI. Domain 2: Organizational Structures for your governance body.

Define what counts as SI, scope, and approval

THE CONTROL SAYS *"Define key concepts in the SI Policy (e.g., definition of SI, scope of SI solution use, SI approval process)."*
In plain terms: Say what you mean: what counts as SI, what is in scope, and how something gets approved.

In practice — how a community health center could implement it

1. Include the five minimum sections

Introduction & Purpose • Scope & Applicability • Relevant Definitions & Terminology • Governance Structure • Responsible SI Lifecycle Management. (We build each one together in the second half.)

2. Start high-level, add detail as you mature

To avoid an overly restrictive policy early, use the minimum sections with flexible descriptors. Take a risk-based approach so staff know which tools need review and how triage works.

3. Reuse what you already have

Fold SI-specific items into existing policies where they fit — recording retention into data management; transparency into consent or disclosure; acceptable use into IT policy.

Note: definitions vary by context

FDA, NIST, WHO, ONC and ISO each define SI-related terms differently. Anchor to the most relevant external standard, acknowledge it may be revised, and add a yes/no decision tree staff can apply to a specific tool.

Example: other policy components

The AMA's Governance for Augmented Intelligence guide adds: risks of SI use, recording retention, transparency guidelines for patients, training requirements — often better placed in existing SOPs. Tampa General Hospital's risk-based process is linked in the playbook.

Reference

Domain 1 Playbook — Control 1.2 guidance and notes; CHAI AI Policy Worksheet with Examples; CHAI Common AI Terms & Definitions; HCA Healthcare Responsible AI Policy (example).

Name the person who keeps it current

THE CONTROL SAYS *"Establish an appropriate individual or individuals with designated oversight for policy changes."*
In plain terms: Name the person — or small group — who owns keeping the policy current.

In practice — how a community health center could implement it

1. Designate an owner

Ideally a senior leader with expertise in digital health, SI, or legal and compliance. In a community health center this is often the Compliance Officer, Medical Director, or an 'SI Coordinator' role added to an existing job.

2. Set the review cadence and triggers

Review at least annually. Also review when new federal or state SI laws pass, or after an adverse event where SI was used. Use language that triggers a required review — not automatic expiration.

3. Document and communicate changes

Version control plus a change log: what changed, why, who approved it, when. Tell the stakeholders who need to know — staff, leadership, the review team.

Note: why this moves fast

Generative SI and LLM-based tools only became common in the last few years, requiring definitions and scope to be updated. The same will likely be true as agentic SI is adopted in healthcare.

Example: two free trackers

Pacific AI Policy Suite (pacific.ai/ai-policies) — quarterly-refreshed policy templates cross-referenced to SI laws and frameworks. Manatt Health AI Policy & Regulatory Tracker — quarterly state and federal summary; its mid-2026 edition counted 280+ state bills introduced in H1 2026, 33 enacted across 22 states.

Reference

Domain 1 Playbook — Control 1.3 implementation guidance and 'Key Tools & Resources'. Potential Challenges table: 'Rapidly Evolving Landscape' and 'Maintenance and Oversight'.

Discussion Questions

Based on the webinar thus far, let's discuss the below... (1 of 3)

1. According to the webinar, what is the most effective way to secure executive buy-in for SI governance?

- A Pitch the newest SI technology and its features
- **B** Lead with organizational risk such as patient safety, regulatory, and reputational risk
- C Promise the lowest possible implementation cost
- D Delegate the decision entirely to the IT department

Explanation. The playbook recommends leading with organizational risk, not a pitch for new technology, framing governance as protecting patient safety, the workforce, and mitigating regulatory, reputational, and financial risk.

2. What is 'shadow SI' as defined in the webinar?

- A SI used only after business hours
- B SI built exclusively by external vendors
- **C** Any SI system, feature, or capability used without going through the organization's established intake, risk assessment, and approval processes
- D SI that has been formally deprecated by the governance committee

Explanation. Shadow SI is any SI system, feature, or capability developed, procured, embedded, or used by staff without going through the organization's established intake, risk assessment, and approval processes.

Discussion Questions

Based on the webinar thus far, let's discuss the below... (2 of 3)

3. From whom should an organization obtain formal approval of the final SI policy?

- **A** Any single department manager
- **B** The highest level of organizational leadership, such as the Board or Executive Committee
- **C** An external consulting firm only
- **D** The vendor supplying the SI tools

Explanation. The playbook directs organizations to obtain approval from the highest level of organizational leadership (e.g., the Board or Executive Committee) and to log that approval.

4. Why does the webinar emphasize a risk-based approach to SI governance?

- **A** It guarantees every SI tool receives the same full review
- **B** It reduces burden on the governance structure and gathers the right information early to avoid future bottlenecks
- **C** It eliminates the need for any executive approval
- **D** It removes the requirement to define SI terms

Explanation. A risk-based process reduces burden on the governance structure and ensures the right information is gathered early in the process to avoid bottlenecks in the future.

Discussion Questions

Based on the webinar thus far, let's discuss the below... (3 of 3)

5. Under Control 1.3, how often should an SI policy be reviewed at minimum?

A Once every five years

B Only when a new CEO is hired

C At minimum annually, with additional triggers such as new laws, new SI capabilities, or new SI risks/adverse events

D Never, once approved

Explanation. Policies should, at minimum, be reviewed annually, with other triggers including new federal or state SI laws or adverse events where SI was used.

STOP 04 • ~30 MINUTES • INTERACTIVE

What Should an SI Policy Include?

A walk through the CHAI AI Policy Worksheet — for every checklist item, what to consider and what it could look like in a community health center

A one-page SI policy to start from

[ORGANIZATION NAME] Super Intelligence (SI) Policy

Draft v0.1 • Policy owner: [name, title] • Approved by: [Board / Executive Committee] • Review date: [annual]

1 Introduction & Purpose

[Organization] is committed to the responsible adoption, deployment, and oversight of AI technologies; this policy establishes the governance framework, accountability structures, and operational standards for SI tools used across the organization, with the goal of ensuring SI use is safe, equitable, effective, and compliant with applicable law.

2 Scope & Applicability

This policy applies to [internally developed / vendor-supplied / system-embedded] SI tools used in [clinical / operational / administrative] settings; workforce members may use publicly available SI tools such as [tools] for [permitted uses] provided no PHI is entered and outputs are verified before use, and the following uses are prohibited: [list].

3 Relevant Definitions & Terminology

For purposes of this policy, Artificial Intelligence means [definition]; Machine Learning means [definition]; Generative SI means [definition]; Clinical Decision Support means software that provides clinicians, patients, or caregivers with knowledge and patient-specific information to enhance health care decisions; [PHI, SaMD, algorithmic bias, model drift, automation bias].

4 Governance Structure

The SI Governance Committee is comprised of [roles/titles] and chaired by [title]; it meets [frequency], reports to [executive/board body] on a [frequency] basis, requires a quorum of [number/percentage], and has [binding / advisory] authority over SI tool approvals.

5 Responsible SI Lifecycle Management

All SI tools shall be classified as [High / Medium / Low] risk based on [criteria]; vendors and development teams must submit [documentation requirements] to the Committee; [high-risk → full review; medium → expedited; low → administrative]; each approved tool shall have a designated owner responsible for ongoing monitoring per the schedule established at approval.

How to use this page

Every line is editable text. Replace each [bracket] with your organization's answer — the five slides that follow walk you through what to consider for each section.

Start small, then grow

Organizations new to governance can include just the minimum sections with high-level, flexible descriptors, and add detail as capabilities mature. This avoids overly restrictive policies early.

Where the sentences come from

Each one is the drafting prompt from the corresponding section of the CHAI AI Policy Worksheet, condensed to a single sentence.

Introduction & Purpose: stating the commitment

Why it matters AI is being adopted rapidly across clinical decision support, documentation, imaging, and operations. Without a foundational policy, organizations risk inconsistent governance, regulatory exposure, and patient safety events.

HCA Healthcare Responsible AI Policy (EC.031)
Ethics & Compliance • effective Oct 1, 2024 • p. 1

PURPOSE: To ensure the responsible development, deployment, and use of Artificial Intelligence ("AI") across the Company following the pillars of the Responsible AI Framework as outlined in the Company's Code of Conduct ("Code"). Responsible AI includes respecting individuals' privacy, promoting transparency, fairness, bias minimization, accountability, and operation in a safe and secure manner that strives to protect individuals from physical, emotional, environmental, and/or digital harm.

Health System SI Oversight Policy

Purpose and Scope

The purpose of this policy is to establish a process for the oversight of artificial intelligence (AI)-enabled solutions that are internally developed or externally procured by the organization and are intended to be used for patient care, including:

- Diagnosis, cure, mitigation, treatment, or prevention of a disease or condition,
- Communication, documentation, or outreach pertaining to individual patients, or
- Clinical operations or planning purposes pertaining to the care of patients.

For CHC's:

- Align responsible SI purpose with CHC mission and guiding principles (not just compliance): underserved groups are most likely to be harmed or excluded by poorly governed SI (or staying behind in innovation).
- Keep it short
- Reference HRSA expectations and FQHC designation where applicable (AI governance can fit within board accountability and QI.)

Introduction & Purpose: stating the commitment

Why it matters SI is being adopted rapidly across clinical decision support, documentation, imaging, and operations. Without a foundational policy, organizations risk inconsistent governance, regulatory exposure, and patient safety events.

- [ ] **States the organization's commitment to responsible SI use**

  What it could look like: '[Health Center] is committed to the responsible adoption, deployment, and oversight of SI technologies' — the worksheet's drafting prompt, with your name in the blank.

- [ ] **Explains why governance is necessary for this organization**

  What it could look like: 'Our EHR vendor is adding SI features and staff are using consumer SI tools; we need one consistent approach to protect patients and PHI.'

- [ ] **Names the relationship to HIPAA, FDA, HRSA, state law, and other policies**

  What it could look like: 'This policy supplements — and does not replace — our HIPAA privacy and security policies and our IT acceptable-use policy.'

- [ ] **States goals: patient safety, equity, compliance, responsible innovation**

  What it could look like: '...with the goal of ensuring SI use is safe, equitable, effective, and compliant with applicable law'

- [ ] **References key risk domains as the rationale**

  What it could look like: One sentence naming patient safety, technology & data, equity, and legal/compliance as the reasons the policy exists.

Scaling for Community Health Centers

A brief, focused introduction (1–2 paragraphs) that references core patient safety and privacy obligations is sufficient.

IN THE CHAT

In one short phrase: the main reason YOUR organization would govern SI.

Reference

AI Policy Worksheet, Section I

Scope & Applicability: who and what is covered

Why it matters Scope determines which AI tools are subject to the policy — and which are not. Policies that define AI too narrowly miss embedded AI inside EHRs. This section also sets the guardrails, including the critical prohibition on entering PHI into public AI tools.

HCA Healthcare Responsible AI Policy (EC.031)

Ethics & Compliance • effective Oct 1, 2024 • p. 1

SCOPE: All HCA Healthcare ("Company") affiliated facilities worldwide including, but not limited to, hospitals, ambulatory surgery centers, home health centers, home health agencies, hospice agencies, physician practices, outpatient imaging centers, urgent care centers, Parallon, joint ventures and all Corporate Departments, Groups, Divisions, and Markets (collectively, "Affiliated Employers" and individually, "Affiliated Employer").

POLICY: This policy applies to: employees', contractors', service providers', and/or vendors' (collectively, "Colleagues") use, engagement, development, or interaction with Company-owned, externally purchased, or publicly available AI Solutions; the data used, stored, and processed for training AI models, and any other tools instrumental in creating Outputs; and AI Solutions in any and all forms, including, without limitation, AI Solutions that are standardized, custom-developed, stand-alone, or bundled with or embedded into any product or service.

1. **Acceptable Use**

   The use of AI Solutions is solely for tasks that contribute directly to Company business objectives and duties and in alignment with the Code, applicable policies, procedures, and law. Colleagues are only allowed to use AI Solutions and tools approved by the Company.

2. **Prohibited Use**

   The Company has identified two specific uses of AI Solutions that are prohibited as a matter of Company policy.

   a. *Dark Patterns*. AI Solutions may not be used to distort, impair, trick, or otherwise interfere with the ability of an individual to make autonomous and informed choices or decisions, or otherwise manipulate a person through subliminal techniques, or so-called dark patterns, to make (or not make) a particular decision or take/refrain from a particular action.

   b. *Exploiting Vulnerabilities*. AI Solutions may not be used to exploit potential vulnerabilities of an individual or to distort or impair their ability to make autonomous and informed choices or decisions or otherwise manipulate or cause physical or psychological harm to themselves or others.

Health System SI Oversight Policy

Purpose and Scope

The purpose of this policy is to establish a process for the oversight of artificial intelligence (AI)-enabled solutions that are internally developed or externally procured by the organization and are intended to be used for patient care, including:

- Diagnosis, cure, mitigation, treatment, or prevention of a disease or condition,
- Communication, documentation, or outreach pertaining to individual patients, or
- Clinical operations or planning purposes pertaining to the care of patients.

This policy does not apply to solutions that do not contain artificial intelligence, such as rules-based algorithms, or to solutions that are not used for patient care. Detailed requirements, including validation and testing, for solutions out of scope of this policy are described in the AI Oversight Standard Operating Procedure (SOP). Questions related to scope and/or applicability of this policy should be directed to the AI Oversight Committee.

Scope & Applicability: who and what is covered

Why it matters Scope determines which AI tools are subject to the policy — and which are not. Policies that define AI too narrowly miss embedded AI inside EHRs. This section also sets the guardrails, including the critical prohibition on entering PHI into public AI tools.

- ☐ **Explicitly covers vendor-supplied AI and AI embedded in existing platforms**

  What it could look like: 'This policy applies to AI features in our EHR and patient portal, contracted AI services, and any AI tool staff use for work — purchased, embedded, or free.'

- ☐ **Names tool categories: internally developed, vendor, system-embedded, clinical, research, operational, administrative**

  What it could look like: A short table with one example of each from your health center: ambient scribe (vendor, clinical); no-show risk score (embedded, operational); translation tool (vendor, administrative).

- ☐ **Defines exclusions (e.g., basic search, spell-check)**

  What it could look like: 'Basic spell-check, grammar suggestions, and standard search functions are not in scope of this policy.'

- ☐ **Lists permitted uses of public AI tools and the conditions**

  What it could look like: 'Staff may use [approved tool] to draft non-clinical communications or summarize public literature, provided no PHI is entered and outputs are reviewed before use.'

- ☐ **Lists prohibited uses**

  What it could look like: No PHI in public or non-contracted tools; no AI as the sole basis for diagnosis or treatment; no unreviewed AI content in records; no unapproved patient-facing tools; no dark patterns or exploiting vulnerabilities.

Note: the consumer AI challenge

Consumer tools are accessible to all staff. Without clear guidance, well-intentioned employees may expose PHI or rely on inaccurate outputs — give practical guardrails while acknowledging real productivity benefits.

Scaling for Community Health Centers

Focus scope on vendor-supplied tools and embedded EHR AI. A simplified permitted / prohibited list is appropriate.

IN THE CHAT

Name one SI tool that is already being used at your org--purchased, embedded, or developed by staff.

Reference

AI Policy Worksheet, Section II (tool categories; permitted and prohibited uses; HCA Healthcare and Mayo Clinic examples).

Definitions & Terminology: build a shared vocabulary

Why it matters Regulatory definitions of SI vary significantly across FDA, ONC, and state regulations. Policies can fail when they define SI too narrowly and miss embedded SI inside the EHR. A shared vocabulary ensures consistent application across the organization.

HCA Healthcare Responsible AI Policy (EC.031)

Ethics & Compliance • effective Oct 1, 2024 • pp. 6–7

AI Solutions or AI Technology means Artificial Intelligence and Machine Learning technologies both individually and collectively, unless otherwise specified within the context of its use.

AI Technology Activity or **AI Technology Activities** means the acquisition, use, development, deployment, modification, distribution, or otherwise making available by or for the Company of AI Technology.

Artificial Intelligence (AI) means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs such as patterns and structures learned from existing data, deep learning, neural networks, and machine learning techniques to, among other actions, (1) perceive real and virtual environments; (2) abstract such perceptions into models through analysis in an automated manner; (3) create new, original content, such as images, text, or music; (4) produce content autonomously that closely resembles human-created output; (5) produce natural language texts based on a given input, such as a prompt, a keyword, or a query and/or (6) use model inference to formulate options for information or action.

Machine Learning means an application of Artificial Intelligence that is characterized by providing systems the ability to automatically learn and improve based on Training Materials or experience, without being explicitly programmed.

Output(s) means any outcome, output, or other result, action, or decision obtained from or otherwise performed by or with the assistance of, an AI Technology and/or AI Solution.

Responsible AI means the area of AI governance that applies across all AI Technology Activities and establishes guidelines to address safety and security, trustworthiness, transparency, fairness, and ethics.

Health System SI Oversight Policy

Acronyms and Definitions

SOP: Standard Operating Procedure

Super Intelligence (SI)-enabled solution(s): For the purposes of this policy, SI-enabled solutions refer to any product, technology, or system that includes one or more algorithm(s) that build relationships between input and output data using machine learning techniques, generative SI, and other statistically-derived models.

Clinical Owner: The Clinical Owner is responsible for representing the end user and use case and serves as a point of contact once the SI-enabled solution is deployed. Activities may include, but are not limited to, providing feedback on solution design according to applicable requirements of this policy, conducting testing, and evaluating safety events and risks to ensure the solution operates as expected before, during, and after deployment.

Model Owner: When an SI-enabled solution is developed internally, the Model Owner is responsible for providing and maintaining information and documentation regarding the technical design, development, and maintenance of the solution, including the model and end user interface of the solution, to ensure it operates as expected and is designed according to applicable requirements of this policy.

Review Checkpoint: A point during the life cycle of the SI-enabled solution where the solution is proactively evaluated by an independent review committee appointed by the SI Oversight Committee to assess whether the model is performing as intended, identify emerging risks and potential biases, confirm adherence to the SI Oversight procedure, and to determine mitigations or improvements, if any, that should be addressed before the next Review Checkpoint.

Definitions & Terminology: build a shared vocabulary

Why it matters Regulatory definitions of AI vary significantly across FDA, ONC, and state regulations. Policies can fail when they define AI too narrowly and miss embedded AI inside the EHR. A shared vocabulary ensures consistent application across the organization.

- ☐ **Defines AI, machine learning and generative AI distinctly**

  What it could look like: AI: 'a machine-based system that... makes predictions, recommendations, or decisions' (FDA). ML: techniques that train AI to improve from data. Generative AI: models that generate new content.

- ☐ **Defines Clinical Decision Support (CDS) and Software as a Medical Device (SaMD)**

  What it could look like: CDS: software that supports or provides recommendations to a clinician, who can review the basis for them (FDA). SaMD: software intended for a medical purpose without being part of a hardware device.

- ☐ **Defines algorithmic bias, model drift and automation bias**

  What it could look like: Automation bias: 'the tendency to over-rely on automated suggestions' (WHO) — pair it with the prohibition on unreviewed AI content in records.

- ☐ **Defines PHI consistently with applicable law**

  What it could look like: Reference the HIPAA definition already in your privacy policy rather than restating it: 'PHI has the meaning given in our HIPAA Privacy Policy.'

- ☐ **Anchors definitions to an external standard**

  What it could look like: 'Definitions are adapted from the FDA Digital Health and AI Glossary and the NIST AI Risk Management Framework and will be updated as those sources change.'

Scaling for Community Health Centers

A simplified glossary of 10–15 essential terms focused on the tools in active use and any terms referenced in policy is sufficient.

Reference

CHAI Common AI Terms & Definitions (FDA / NIST / WHO / ONC / ISO); AI Policy Worksheet, Section III; Playbook note 'Definitions Vary by Context of Use'.

Key terms to define (examples)

AI • Machine Learning • Generative AI • Clinical Decision Support (CDS) • Software as a Medical Device (SaMD) • Algorithmic Bias • Model Drift • Automation Bias • Protected Health Information (PHI)

Governance Structure: who decides, and how

Why it matters Effective governance requires clear structures, defined roles, and appropriate resourcing, matched to organizational complexity. Define who sits on the governance group, how decisions are made and reported upward, and how often it meets.

HCA Healthcare Responsible AI Policy (EC.031)

Ethics & Compliance • effective Oct 1, 2024 • pp. 1, 7

3. **Responsible AI Governance Council**

   The Responsible AI ("RAI") Governance Council has been established to address areas of stakeholder engagement, ethical considerations, policy development, risk management, and compliance aspects pertaining to AI Solutions.

Responsible AI (RAI) Governance Council means the Sponsors, Steering Committee, and Advisory Committee as confirmed from time-to-time within the Company who provide sponsorship of the Responsible AI Program and focus on impacts to the enterprise, funding, timeline, and major risks arising from AI Technology Activities.

Director of Responsible AI means the Colleague(s) identified by the RAI Governance Council from time to time who is responsible for the management and oversight of the Responsible AI Program.

Health System SI Oversight Policy

SI Oversight Committee

Oversight of SI-enabled solutions is performed by the SI Oversight Committee.

The SI Oversight Committee is comprised of members from the health system and its affiliated university, including its schools of medicine and nursing, to ensure that algorithmic solutions used in the delivery of healthcare at the health system are safe, fair, equitable, and of high quality.

Clinical Owner: The Clinical Owner is responsible for representing the end user and use case and serves as a point of contact once the SI-enabled solution is deployed. Activities may include, but are not limited to, providing feedback on solution design according to applicable requirements of this policy, conducting testing, and evaluating safety events and risks to ensure the solution operates as expected before, during, and after deployment.

Model Owner: When an SI-enabled solution is developed internally, the Model Owner is responsible for providing and maintaining information and documentation regarding the technical design, development, and maintenance of the solution, including the model and end user interface of the solution, to ensure it operates as expected and is designed according to applicable requirements of this policy.

Governance Structure: who decides, and how

Why it matters Effective governance requires clear structures, defined roles, and appropriate resourcing, matched to organizational complexity. Define who sits on the governance group, how decisions are made and reported upward, and how often it meets.

- ☐ **Names who is represented**

  What it could look like: Medical Director (chair), Compliance / Privacy Officer, IT lead, Operations lead, user group representative, and a patient advisory council member.

- ☐ **Defines roles and accountabilities**

  What it could look like: 'The SI Coordinator runs intake and keeps the SI inventory; the Medical Director signs off on clinical tools; the Compliance Officer reviews PHI handling and the BAA.' Include clinical/user champions to help with decision-making and buy-in

- ☐ **States how and how often the group reports to leadership or the board**

  What it could look like: 'The SI Coordinator reports to the Executive Team quarterly and to the Board annually, including the inventory of approved tools.'

- ☐ **States meeting cadence, quorum, and how urgent decisions are made**

  What it could look like: 'The review team meets monthly or as needed; between meetings the Medical Director and IT lead may jointly suspend a tool.'

- ☐ **States whether the group recommends or has binding authority**

  What it could look like: 'The review team recommends; the CEO (or Executive Team) approves.' — or give the team binding authority for low-risk tools.

Scaling for Community Health Centers

An SI Coordinator plus an ad-hoc review team, with external consulting support as needed. A lean structure with clear role definitions is sufficient.

IN THE CHAT

Who would be on your review team? Type the roles — not names.

Reference

AI Policy Worksheet, Section IV; Domain 2: Organizational Structures Playbook (next in this series).

Lifecycle Management: intake, review, and monitoring

Why it matters Responsible AI use requires governance across the full lifecycle — intake and risk classification through approval, deployment, and monitoring. This section sets the risk-based approval process and what vendors and internal teams must provide.

HCA Healthcare Responsible AI Policy (EC.031)
Ethics & Compliance • effective Oct 1, 2024 • pp. 2–3

PROCEDURE:

1. **Seeking Use of AI Solution**
    a. Those seeking an external solution leveraging AI must follow Company solution submission and approval processes. A RAI Questionnaire will be submitted as a function of the AI Solution submission and approval processes and RAI risk assessment performed.
    b. Those seeking to engage in development, co-development, or deployment of solutions (AI Development Lifecycle) leveraging AI must first contact the Director of Responsible AI prior to initiating any development or deployment.
        i. A Use Case Risk Assessment will be performed during solution ideation.
        ii. A Data Risk Assessment will be performed prior to solution development.
        iii. A Validation Risk Assessment will occur prior to moving to solution production or deployment.
    c. All solutions leveraging AI will be subject to periodic audit for the duration of the solution lifespan until deprecation to ensure compliance and alignment with the RAI framework. Any solution leveraging AI that undergoes an internally or externally initiated or offered upgrade or new version shall be subject to an audit to ensure continued compliance and alignment with the RAI.
    d. Solution approval and review is subject to a risk score and the RAI Governance Council reserves the right to review or rescind approval at any time. Solutions proposed for review before the RAI Governance Council and Steering Committee will be prioritized on risk versus enterprise impact.
        i. Low Risk Score: Automatic approval and awareness of the RAI Governance Council.
        ii. Medium Risk Score: Review by steering committee, mitigation recommendations, and awareness of RAI Governance Council.
        iii. High Risk Score: Review by steering committee, mitigation recommendations, referral to full RAI Governance Council, and issuance of a written go/no-go decision by the RAI Governance Council.
        iv. Critical Risk Score: Review by steering committee, mitigation recommendations, referral to full RAI Governance Council, and issuance of a written go/no-go decision by the RAI Governance Council.
    e. The Company will maintain a process for creation and maintenance of an AI solution inventory. For further information, contact the Director of Responsible AI.

What to notice

- Risk-based: each solution gets a risk score (low to critical) that sets the level of review and approval.
- Lifecycle checks: use case, data, and validation risk assessments come before development and production.
- Ongoing oversight: periodic audits, the right to rescind approval, and a maintained AI solution inventory.

Lifecycle Management: intake, review, and monitoring

Why it matters Responsible SI use requires governance across the full lifecycle — intake and risk classification through approval, deployment, and monitoring. This section sets the risk-based approval process and what vendors and internal teams must provide.

Academic Health System AI Oversight Policy
Governance & oversight committee • health system anonymized • p. 4

SI Oversight Process – Evaluation

Evaluation requirements for SI-enabled solutions are commensurate with risk.

The SI Oversight Committee triages registered SI-enabled solutions based on risk and determines the applicable review process and evaluation requirements.

Risk categories and evaluation requirements are defined and maintained by the SI Oversight Committee and documented in the SI Oversight SOP.

These categories include:

- Out of scope for SI Oversight Committee review (low risk),
- Fast-Track Review (moderate risk), or
- Full Review (high risk)

SI-enabled solutions that are determined to be in scope of this policy must satisfy the applicable testing and evaluation requirements described in the SI Oversight SOP before, during, and after deployment, and throughout use.

When SI-enabled solutions are out of scope for SI Oversight Committee review, Model Owners and/or Clinical Owners are responsible for performing independent validation and testing of the SI-enabled solution for patients.

SI Oversight Process – Review Checkpoints

The SI-enabled solutions that are in scope for SI Oversight Committee review are required to pass Review Checkpoints before, during, and after deployment, and throughout use.

The SI Oversight Committee, or a review committee that is appointed by it, determines whether SI-enabled solutions satisfy the Guiding Principles at Review Checkpoints before, during, and after deployment, and throughout use in patient care. At each Review Checkpoint, the assigned committee independently evaluates whether the model is performing as intended, identifies benefits and any emerging risks, confirms adherence to the Guiding Principles, and determines mitigations or improvements, if any, that should be addressed before the next Review Checkpoint.

Safety Events

Safety Events that may involve SI-enabled solutions must be reported to the SI Oversight Committee and follow reporting requirements of the organization's Safety Reporting Policy.

Lifecycle Management: intake, review, and monitoring

Why it matters Responsible SI use requires governance across the full lifecycle — intake and risk classification through approval, deployment, and monitoring. This section sets the risk-based approval process and what vendors and internal teams must provide.

- ☐ **Classifies tools by risk before review**

  What it could look like: Two tiers: 'Standard review' for anything touching patient care or PHI; 'Expedited review' for administrative tools with no PHI. Grow to three tiers later.

- ☐ **Describes the workflow: intake → classification → review → decision → conditions**

  What it could look like: A one-page intake form → the SI Coordinator assigns the tier → review team → CEO decision → conditions (training, monitoring, limits) recorded in the SI inventory.

- ☐ **Lists what vendors must provide at intake**

  What it could look like: Intended use; performance in populations like ours; how PHI is handled and retained; bias testing; FDA / ONC status; BAA plus SOC 2 or HITRUST.

- ☐ **Names a designated owner and monitoring cadence for each tool**

  What it could look like: 'Each approved tool has a named owner who reviews the vendor's performance report quarterly and reports exceptions to the review team.'

- ☐ **Defines incident reporting and decommissioning**

  What it could look like: SI-related near-misses go through the existing incident-reporting system; the owner or Medical Director can suspend a tool pending review.

Scaling for Community Health Centers

A simple two-tier (standard / expedited) review may be sufficient. Rely on vendor-supplied monitoring reports — and specify what they must contain and how often.

IN THE CHAT

Standard or expedited? (1) an ambient scribe (2) a scheduling optimizer. Type your answers.

Reference

AI Policy Worksheet, Section V; Playbook note 'Risk-Based Approach' (Tampa General example); Subdomains 4.1 and 4.4 (later in this series).

Approve, share, and maintain

§4 Approve, share, and maintain

Secure executive commitment

Artifact: Written commitment and an initial budget or repurposed resources

Obtain formal approval

Artifact: Cross-functional review; Board or Executive Committee approval, logged

Make it accessible; inform staff

Artifact: Published policy; champion; attestation; usage request form

Maintain it (Control 1.3)

Artifact: Named owner; annual review and triggers; change log; communicate

§5 A suggested first 30 days

Week 1	Identify your executive sponsor and name a policy owner. Frame the case around risk. Pick your 3–5 person drafting team.
Weeks 2–3	Block 90 minutes. Use Policy Worksheet to guide team through each section; fill the drafting prompts; keep it high-level; borrow definitions from the CHAI guide where appropriate.
Week 4	Cross-functional review; check against existing policies and state/federal rules; route to the Board or executive committee; log the approval.
Then	Publish; brief staff; attach a usage request form; set the annual review date and the triggers that move it up.

In the chat: what's the one thing you'll do in the next two weeks?

Take these with you

Materials from today

Domain 1: Policy Playbook
https://www.chai.org/workgroup/cross-cutting/ai-governance

SI Policy Worksheet with Examples
Five sections, formulation questions, drafting prompts, scaling tips
Policy Worksheet

Common SI Terms & Definitions
Harmonized FDA / NIST / WHO / ONC / ISO definitions for your glossary
Common Definitions

Pacific AI Policy Suite • Manatt Health AI Policy Tracker
Free resources for tracking regulatory change

Example policies
HCA Healthcare Responsible AI Policy
Tampa General Risk-based Review Process

Next in the series

01	**AI Policy**	DONE
02	Organizational Structures	
03	Organizational Resources	
4.1	Responsible AI Lifecycle Management	
4.2	Risk & Impact Assessments	
4.3	Responsible Data Management	
4.4	Third-Party Management	
4.5	Education, Training & Feedback	

Domain 4 playbooks are organized as five subdomains.

THANK YOU

Questions?

Coalition for Health AI • merage@chai.org • anthony@chai.org

The CHAI AI Governance Playbooks are voluntary, educational resources. They do not establish a standard of care and are not legal advice. Examples and tools named are illustrative and do not indicate endorsement by CHAI. Playbook content licensed CC BY-NC-ND 4.0.

Workshop Evaluation Survey

SCAN QR CODE

Please share your feedback on this session. This should take less than 3 minutes to complete.

Survey Link:
https://online2.snapsurveys.com/interview/b79cbaea-263f-4e11-bd53-63dfd472fe02

