Domain 1: AI Policy Playbook

Guiding responsible AI adoption

Before You Start

For more about the CHAI AI Governance Framework, methods used to develop these playbooks, and information on how to use them, please refer to the CHAI AI Governance Playbooks: Introduction & Background document. Note that **baseline operational controls** were consensus-defined by Healthcare Delivery Organizations (HDOs) as "must-haves" for health AI governance. The **implementation guidance** that accompanies the baseline operational controls are **suggestions/recommendations** based on known practices. It is assumed and appropriate that implementation will differ based on organizational goals, resources, and maturity.

Disclaimer

The implementation guidance, frameworks, and examples contained herein represent generalized practices and considerations drawn from across the field of responsible AI deployment in healthcare. They are offered as one set of perspectives to inform your thinking not as the "right way" or "only way" to build, govern, or operate an AI program.

Every organization operates within its own unique context: differing patient populations, clinical priorities, regulatory environments, resource constraints, technical infrastructure, cultural norms, and stage of AI maturity. While the specific methods, timelines, and depth of implementation will appropriately vary across organizations, the underlying controls, principles, risks, and considerations raised throughout this playbook are intended to be engaged with thoughtfully and in good faith by any organization deploying AI in healthcare.

Organizations should prioritize state and federal regulatory requirements, especially for AI solutions that meet the definition of a Software as a Medical Device (SaMD), as defined by the FDA.

Not a Standard of Care. This playbook does not establish, define, or reflect a legal, regulatory, professional, or industry standard of care, nor is it intended to serve as one. It is a voluntary, aspirational, and educational resource designed to support organizational learning and dialogue. The field of healthcare AI is rapidly evolving, and reasonable organizations and professionals may differ on how to approach the issues discussed herein. The practices described may not be appropriate, feasible, or advisable for every organization, every use case, or every point in time. Adoption of any particular practice described in this playbook is entirely voluntary, and a decision to adopt, adapt, partially implement, or decline to implement any practice should not be construed as evidence of compliance with, or deviation from, any standard of care, duty, or legal obligation. This playbook is not intended to be used, and should not be used, as a benchmark, checklist, or measuring stick against which the conduct of any

organization, clinician, or other party is evaluated in any litigation, regulatory proceeding, accreditation review, or similar context.

General Informational Purpose. The information in this playbook is provided for general informational and internal operational purposes only. It is intended to support organizations in thinking through the design and stewardship of their AI programs and should be used in conjunction with from clinical leaders, technical experts, compliance officers, ethicists, patient and community representatives, and other advisors familiar with the specifics of your organization.

Not Legal Advice. Nothing in this playbook constitutes legal advice, and it should not be relied upon as a substitute for advice from a qualified attorney licensed to practice in the relevant jurisdiction. Nothing in this document creates an attorney-client relationship between the reader and CHAI, its officers, directors, employees, or contributors. The contents reflect general information and observed practices as of the date of publication and may not account for changes in law, regulation, technology, or circumstance that occur thereafter.

No Warranty. While reasonable efforts have been made to ensure accuracy, no representation or warranty is made as to the completeness, timeliness, accuracy, or suitability of the information for any particular purpose. Laws, regulations, and best practices vary by jurisdiction and evolve over time, and the application of those laws and practices to specific facts can lead to different outcomes.

Illustrative Examples. Any tools, resources, examples, or technology product described or presented are for educational and illustrative purposes only and do not represent endorsement by CHAI or guarantee specific results. Readers should consult qualified legal counsel and relevant experts for specific terms, products, or services.

Independent Judgment Required. Readers should consult with qualified legal counsel, clinical experts, and other appropriate professional advisors before taking, or refraining from taking, any action based on the information in this playbook. The decision to adopt, modify, or disregard any practice described herein rests solely with the reader and their organization. CHAI expressly disclaims any and all liability for any loss, damage, or other consequence, including but not limited to claims of negligence, breach of duty, or failure to meet any standard, arising directly or indirectly from reliance on, use of, citation to, or reference to the contents of this document.

At a Glance

		
Formalize Documentation & Secure Executive Buy-in	**Define Key AI Concepts & Scope**	**Assign Oversight & Maintain Policy**

This **Domain 1: AI Policy** Playbook is designed to guide HDOs through the steps of developing a formal, operational AI policy that establishes clear accountability and defines the scope of AI governance within the organization. It offers actionable guidelines, real-world examples, and essential tools to navigate the complexities of modern governance.

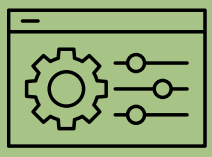

Baseline Operational Controls for Domain 1: AI Policy

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

AI Policy: The Foundation

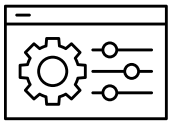

Control 1.1: Develop a formal AI Policy document, approved by organizational leadership, and make accessible to all relevant staff.

A formal, organization-wide policy standardizes the ethical and safe development (when applicable), evaluation, implementation, use, and monitoring of AI solutions.

Implementation Guidance

Task 1: Secure Executive Commitment

Policy development requires **executive sponsorship** to ensure compliance and resource allocation. If necessary, address the challenge of change management in leadership by preparing a business case focused on risk mitigation, protecting workforce, and patient safety, not just technology.

Action: Schedule strategic planning sessions and one-on-one meetings with the CEO, CMO, CNO, CNE, or CIO (or related roles) to secure written commitment and an initial budget (or resource repurposing) for governance. This commitment is necessary for subsequent executive approval.

Note: Prioritization for Securing Executive Buy-in

Executive investment in establishing an AI governance process and program can have several challenges. This can be due to resource constraints, low AI literacy/trust, different motivating factors, or general change management concerns. Given the multi-stakeholder nature of effective AI governance, it may be necessary to gain buy-in from executives across multiple departments (e.g. CEO, Innovation, IT, Operation, Clinical, etc). Ideally approaches should be personalized to the executive role and their priorities.

Steps that have helped organizations gain executive buy-in for AI governance include:

- **Leading with organizational risk, not pitch for new technology.**

  Governance enables:

  - Patient safety protection
  - Workforce protection
  - Regulatory and liability risk mitigation
  - Reputational risk prevention
  - Financial risk prevention
  - Bias mitigation
- **Using real failure stories and near-misses**
  - Examples of AI failure, misuse, or unintended consequences (internally or externally) are highly effective in making the risk tangible and urgent.
- **Demonstrating how governance can lead to faster and safer adoption, not slower**
  - Executives are more supportive when governance is framed as:
    - Enabling triage of low-risk solutions
    - Speeding up safe approvals
    - Preventing bottlenecks caused by ad-hoc reviews
- **Showing how governance protects clinicians and staff**
  - Governance structures can help to clarify responsibility and protect clinicians, their patients, and staff from unsafe AI solutions or unsafe use of AI solutions
  - Workforce impacts (burnout, trust in AI solutions, workflow disruption, and accountability burden on clinicians) were repeatedly cited as a strong executive motivator.
- **Providing clear Return on Investment (ROI) and Value on Investment (VOI) framing that is tied to risk and sustainability**
  - ROI was consistently described as important, but expanded to include VOI that went beyond financial return alone.
  - Stronger framing focused on:
    - Avoiding costly failures
    - Reducing rework and failed pilots
    - Protecting investments by stopping poor AI solutions early
    - Improving prioritization of limited capital
    - Above stated workforce impacts

Task 2: Obtain formal approval

The draft policy should undergo **cross-functional review** with legal/compliance, IT, ethics, operations, and clinical teams to ensure it does not conflict with existing organizational policies (e.g., cybersecurity or data governance), and existing state, local, or federal requirements (see Control 2 implementation guidance for more on what is important to include in a p. Obtain approval from the highest level of organizational leadership (e.g.,

the Board or Executive Committee). Consider engaging your defined governance body to help with the development and approval of the policy (see Domain 2: Organizational Structures Playbook)

Artifact: Obtain executive leadership approval for the final AI policy document and log this approval.

Task 3: Ensure Accessibility & Inform Staff

Once approved, the policy should be disseminated and accessible to all relevant staff. Communication should effectively convey the value of the policy.

Action: Publish the final AI policy to all staff via internal portals, the organizational website, and incorporate key sections into role-specific training materials for relevant staff. Note that in addition to establishing a policy, it may be important to develop processes for its implementation, adoption, and enforcement, as appropriate. This may include things like adding incentives and safeguards to ensure policy is being followed, developing methods for identifying breaches, hosting staff town halls to answer questions/concerns, and embedding AI policy training into standing training processes.

Note: Getting In Front of Shadow AI

When it comes to disseminating AI policies, repetition and approachable forums to share questions and concerns are key.

Some organizations assign an appropriate champion to take the document/policy to department and staff town halls and meetings to inform staff. Others included a policy attestation as part of their dissemination process to ensure that relevant staff had seen the policy and understood its implications.

This dissemination process may also be a good opportunity to attach a "usage request form" to identify things like "shadow AI".

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

E.g.

- A department uses a public generative AI tool to summarize patient notes.
- A vendor turns on an AI feature inside the EHR or another system without a separate governance review.
- A team builds a small prediction or triage model locally using operational or clinical data.
- Staff upload documents or data into third-party AI solutions outside approved systems.
- Use of copilot or direct to consumer AI products to assist with tasks at work.

Adding a usage request form, can help organizations identify:

a) which AI solutions are currently being used by staff--this provides a starting place for testing out the AI policy and governance process

b) which AI solutions are being requested by staff

Other complementary methods for identifying shadow AI may include network-level and data loss prevention controls that can be used to identify, monitor, and block the unauthorized sharing, transfer, or use of sensitive information.

Note that when AI solutions are included as part of an update/embedded product within an existing vendor platform (e.g. EHR), it may be important to update or add an addendum to existing contractual terms or to internal standard operating procedures to ensure appropriate notification and governance due diligence (see Domain 4.4: Third Party Management Playbook for more on this)

AI Policy: Define Key Concepts & Scope

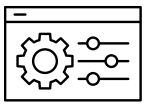

Control 1.2: Define key concepts in the AI Policy (e.g., definition of AI, scope of AI solution use, AI approval process).

Implementation Guidance

At minimum, an AI policy should include:

1. Introduction & Purpose
2. Scope & Applicability: Description of when/which tools require governance (internally developed, vendor developed, system embedded, clinical, research, operational, administrative etc.) and when they do not, as well as a description of permitted and prohibited uses (e.g. prohibition of use of PHI in public AI tools).
3. Relevant Definitions & Terminology
4. Governance Structure: Who is involved in the governance committee, how is it structured, how does reporting occur, and how often do they meet? (See Domain 2: Organizational Structures Playbook for more information on which roles and responsibilities are important to include here).
5. Responsible AI Lifecycle Management: A description of the risk-based governance processes and procedures for approval across the AI lifecycle.
    i. Including intake and monitoring documentation requirements for vendors/development teams (please see Subdomain 4.1: AI Lifecycle Management and Subdomain 4.4: Third Party Management Playbooks for more information.)

In order to avoid developing overly restrictive AI policies early, organizations new to their AI governance journeys might consider including just the minimum policy sections with high level, flexible descriptors, and updating with greater detail as their organization and capabilities mature. Here is an example of how organizations may start out and then expand their AI policies.

Note: Risk-Based Approach

It is recommended that organizations take a risk-based approach to governance, identifying what kinds of solutions are in or out of scope of formal governance, and the rigor of governance required depending on the risk of the AI-solution.

The "Scope" and "Governance Structure" sections of the policy should help individuals know which solutions will need to go through a governance/approval process, and if governance applies, how triage of solutions will go. There are different ways that organizations approach this, but the benefit of a risk-based process is that it reduces burden on the governance structure and ensures that the right information is gathered early in the process to avoid bottlenecks in the future.

See here for an example risk-based process from Tampa General Hospital

Other sources, such as the AMA's Governance for Augmented Intelligence Guide include policy components such as:

- Risks of AI use
- Recording retention policy
- Transparency guidelines for patients/users
- Training requirements

Some of these additions may be more use case specific and variable, and may subsequently be more appropriate as updates in new or existing policies and standard operating procedures. For example, recording retention information may be part of a data management policy or standard operating procedure, and transparency may be a part of existing consent or disclosure policy.

Where possible, update existing processes, policies, and procedures to minimize resource strain.

The same is true for governance structures. It may not be feasible to include an entirely separate committee responsible for intake, evaluation, and monitoring of AI solutions, but identifying which roles might be required for different parts of AI lifecycle management, is important.

Note: Definitions Vary by Context of Use

Even across regulatory and standard frameworks, there is considerable variability in how the term "Artificial Intelligence" is defined. This is in part due to differences in their focus (e.g. regulatory, risk management, ethics, health IT certification, etc).

For example:

FDA (2024): 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 to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action. (Source: FDA Digital Health and Artificial Intelligence Glossary-Educational Resource)

NIST AI RMF (2023): Engineered or machine-based systems that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. (Source: NIST AI RMF)

World Health Organization (2021): Technology designed to perform tasks typically requiring human intelligence, including learning, reasoning, problem solving, and adapting to new situations. Emphasizes ethical deployment in healthcare. (Source: WHO Ethics and Governance for Artificial Intelligence for Health)

Office of the National Coordinator for Health IT (2024): Technology 'supplied by' certified health IT vendors that use computational techniques to analyze data and make predictions or recommendations.

International Organization for Standardization (ISO 22989, 2022): Research and development of mechanisms and applications of engineered system that generates outputs such as content, forecasts, recommendations, or decisions for a given set of human-defined objectives. Source: Information Technology--Artificial Intelligence--Artificial Intelligence Concepts and Terms.

See here for a list of commonly used terms and definitions.

Ingestion Modalities of AI Solutions for Governance Policy Consideration

AI solutions enter organizations through different ingestion modalities and governance requirements should reflect those differences. Refer to the **Background Context for Subdomain 4.1: Responsible AI Lifecycle Management** Playbook for more information about each ingestion modality. Note that in alignment with a risk-based approach, organizations may want to differentiate policy considerations across different types of AI (e.g. traditional machine learning, generative AI, agentic AI, etc) or different levels of decision autonomy (e.g. fully autonomous agentic solutions that make decisions, versus those that make recommendations that a human ultimately decides on.)

Recommended Reading: Health Systems Only Govern the Tip of the Iceberg

Key Tools & Resources

CHAI Tools and Resources

- Policy Worksheet with Examples
- CHAI Partnered Governance Platforms
- Illustrative Examples of how policies can evolve across maturity

Other Resources to help with AI Policy Development

- AMA Governance Playbook (Step 4 Policy)
- AI Career Pro Policy Development Resources by James Kavanagh
- Developing AI Policies for Public Health Organizations: A Template & Guidance

Example Responsible AI Policy

- HCA Healthcare: Responsible AI Policy

AI Policy: Changes & Maintenance

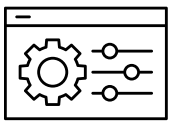

Control 1.3: Establish an appropriate individual or individuals with designated oversight for policy changes.

Policies without a named owner tend to be updated informally, inconsistently, or not at all. Control 1.3 addresses this by ensuring that responsibility for maintaining the AI policy is explicitly assigned, not assumed. A designated policy owner creates a clear point of accountability for initiating reviews, incorporating regulatory or operational changes, and ensuring the policy remains aligned with the organization's evolving AI portfolio. Without this role, even well-written policies can become outdated.

Implementation Guidance

The rapidly evolving landscape of AI regulations necessitates review and maintenance of AI and related policies. Designated oversight ensures timely updates are triggered by changes in laws, new AI adverse events, or new/emerging technologies. For example, while traditional machine learning technologies have been around for many years, generative AI and LLM based technologies only became more common in the last few years, requiring definitions and scope to be updated. The same will likely be true as agentic AI becomes more widely adopted in healthcare.

1. **Designate someone to oversee AI policy changes**: This individual should be a senior leader ideally, with expertise in digital health, AI or legal and compliance.
2. **Establish a regular review cadence:** Define what triggers a policy review and update. Policies should, at minimum, be reviewed annually. Other triggers for review may include when new federal or state AI laws are passed, or adverse events arise where AI was used.

3. **Document changes:** Use version control and log what was changed, why the change was made, who approved it and when, for accountability and auditability.
4. **Communicate changes** to necessary stakeholders.

Key Tools & Resources

Resources for tracking regulatory policies and updates:

- Pacific AI Policies Suite (healthcare inclusive, but not exclusive)
- Manatt Health AI Policy & Regulatory Tracker

Potential Challenges in Developing & Maintaining an AI Policy

Below are a list of challenges that organizations said they faced with developing, implementing, and maintaining AI policies.

Challenge/Gap	Details	Mitigation
Rapidly Evolving Landscape	Difficulty in keeping policies current due to fast-moving regulatory and product changes.	Use a modular policy structure that separates stable principles from technology-specific guidance, so only targeted sections need updating. Assign a policy owner to monitor regulatory changes on a defined cadence (e.g., semi-annual review).
Accessibility and Communication	Challenge of making policies accessible and communicating their value effectively.	Create tiered policy materials: a full governance document, a plain-language staff summary, and role-specific one-pagers for clinical and executive audiences. Pair policy with decision trees that help staff determine whether and how the policy applies to a specific tool. Designate departmental AI liaisons to serve as translators and maintain a recurring communication cadence.
Leadership Engagement and Literacy	Low AI literacy in leadership hinders endorsement and commitment.	Establish a recurring AI literacy program for board and C-suite (e.g., annual education day, quarterly committee updates). Frame governance in terms leadership already owns: patient safety risk, regulatory liability, and reputational exposure. Use real-world AI failure case studies as evidence and connect governance metrics to existing organizational performance dashboards.
Scope and Granularity	Balancing scope (all AI vs genAI) and granularity (high-level vs detailed).	Build a tiered structure: one universal high-level policy covering all AI, with supplemental annexes for higher-risk categories (e.g., GenAI, clinical decision support). Use risk stratification to calibrate documentation and review burden to the risk level of the tool — avoiding full review requirements for low-risk administrative solutions. Pair with a scope decision tree to reduce interpretive ambiguity.
Alignment with Existing Policies	Ensuring AI Policy does not conflict with cybersecurity, safety, quality, or data governance policies.	Conduct a policy alignment audit before finalizing, mapping AI policy provisions against existing cybersecurity, data governance, quality oversight, and patient safety policies. Involve cross-functional stakeholders (IT security, legal, compliance, clinical informatics, quality, etc.) in drafting. Use explicit cross-referencing language to defer to existing domain policies rather than restating or contradicting them, and assign a policy coordinator to maintain alignment over time.

Defining AI Concepts	Lack of consensus on formal and operational definitions of AI.	Anchor definitions to an established external standard (e.g., NIST AI RMF, ISO/IEC 22989:2022, FDA definitions, academic/industry literature) and acknowledge that definitions are subject to revision. Supplement formal definitions with a practical decision tree translating the definition into yes/no questions staff can apply to a specific tool. Include illustrative examples of what is and is not in scope to reduce interpretive variance.
Maintenance and Oversight	Burden of frequent review, re-evaluation, oversight, sunset clauses, and change logs.	Assign named policy ownership with a documented review calendar. Use a tiered review schedule (lightweight annual check vs. comprehensive biennial review). Maintain a standardized change log within the document. Use language that triggers a required review — not automatic expiration — to prevent the policy from becoming outdated without deliberate action.

Next Steps

Links below will connect you to other AI governance playbooks in this series.

Domain 1: AI Policy

Domain 2: Organizational Structures

Domain 3: Organizational Resources

Domain 4: Organizational Processes

- Subdomain 4.1: Responsible AI Lifecycle Management and Use
- Subdomain 4.2: Risk & Impact Assessments
- Subdomain 4.3: Responsible Data Management and Use
- Subdomain 4.4: Third Party Management
- Subdomain 4.5: Education, Training, and Feedback

Attribution

CHAI AI Governance Playbooks were developed through the contributions and collaboration of more than 150 organizations across the healthcare ecosystem. CHAI extends its sincere appreciation to all participating organizations and contributors for their time, expertise, and thoughtful feedback throughout the development process. We would also like to provide special recognition to the individuals below who requested attribution by name for their contributions to these efforts:

- Josh Hjelmstad, MHA, AVIA Health
- Sandra Gregg, DNP, MHA, RN, PMP, LSSBB, CCMP, HCD, Booz Allen Hamilton, Inc.
- Vidhi Vinay Kothari, Carnegie Mellon University
- Stephen McLaughlin, PhD, Children's National Hospital
- Ryan Marshall Felder, PhD HEC-C, Cleveland Clinic
- Po-Hao Chen, MD, MBA, Cleveland Clinic Foundation
- Katherine Mulready, JD MPS, cliexa
- Holden Groves, MD, Columbia University/NewYork-Presbyterian
- Shakira J. Grant, MBBS, MSCR, CROSS Global Research & Strategy
- Namrata Rastogi, MBBS MRCGP MS, Enigma Ventures
- Aleocidio Sette Balzanelo, MD, MBA, FOLKS
- Rocco Orlando, MD, Hartford HealthCare
- Michael Blumental, HealthLeap
- Romanus M. Joseph, Innovative Health Accountable Care Organization (ACO)
- Rebecca G. Mishuris, MD, MS, MPH, Mass General Brigham
- Nasibeh Zanjirani Farahani, PhD, CSM, Mayo Clinic
- Jason Gillman, MD, FIDSA, MCG Health
- Douglas Graham, Individual Contributor
- Nicole Petroski, Mercy Health Services
- Haris Shuaib, Newton's Tree
- Rémy Ibarcq, Individual Contributor
- Srikar Nallan, MS, Stanford Healthcare
- Jaimie Weber, MD, Tampa General Hospital
- R. Brandon Hunter, MD, FAAP, Texas Children's Hospital
- Angela L. Lipscomb, JD, LL.M, The University of Texas MD Anderson Cancer Center
- Christopher H. Lee, MBA, BSN, RN-BC, UCLA Health
- Abby Steele, Individual Contributor
- Shiba Kuanar, University of Minnesota
- Joshua Miller, JD, CHC, University of Rochester Medicine
- Anne Travis, MD, MSc, Wolters Kluwer
- Michael L. Shaw, JD, ZS

License