Domain 2: Organizational Structures 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

This playbook outlines the steps and processes for establishing an organizational structure for AI governance.

Background Context for Domain 2: Organizational Structure

As organizations think about how to utilize existing and develop new organizational structures to support health AI governance, it is important to be aware of the resources available and needed, and to define who will be responsible and accountable for each of the functions necessary.The resources and staff for health AI governance may be found both within and outside your organization. Below is a summary of the functions that need responsibility and accountability.

Risk, Bias, and Impact Assessments

A standard risk, bias/fairness, and impact evaluation, and a risk-based governance approach supports safe, consistent, and accountable AI governance (See Subdomain 4.2: Risk & Impact Assessments Playbook). A tiered risk categorization process, and a tiered governance process can help reduce burden on the organization, and speed up review and implementation where risks are low and benefits high.

Responsible AI Lifecycle Management & Use

Organizations should operationalize Responsible AI principles across the AI lifecycle (an example displayed below; See Subdomain 4.1: Responsible AI Lifecycle Management & Use Playbook). Depending on whether your organization develops AI solutions internally or not, stages 2-3 may not be relevant, but stages 1 and 4-6 will likely apply to all organizations.

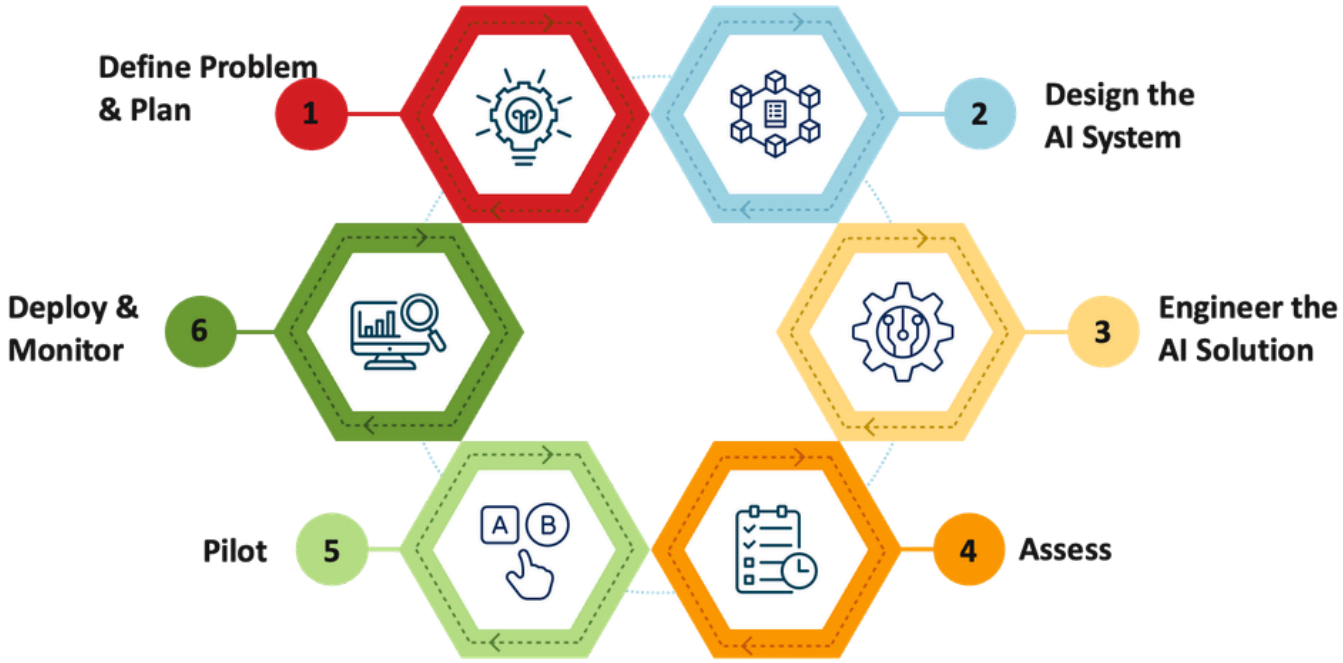

1

- Engage stakeholders to define the problem & perform root-cause analysis
- Identify solution & plan future state
- Gather business requirements
- Assess feasibility, potential for impact, & prioritization
- Make procure/ build/partner decision

2

- Select/understand model task & architecture
- Capture design & technical requirements or determine best solution to meet business requirements
- Design solution application & system workflow according to human-centered design principles
- Design deployment strategy with end users
- Design risk management, monitoring & reporting plan

3

- Access data
- Prepare data
- Develop data management plan
- Train & tune model

4

- Conduct installation qualification (when applicable)
- Validate local system performance (when applicable)
- Execute prospective, silent evaluation
- Establish risk management plan
- Train end users
- Test usefulness
- Ensure compliance with applicable healthcare regulations & standards

5

- Implement small-scale pilot to assess real-world impact
- Execute and update risk management plan
- Educate & train users on AI application reporting
- Assess usefulness and adoption

6

- Deploy at a larger scale on a general population
- Audit AI system to inform whether to maintain, refine or sunset
- Conduct ongoing risk management

While there are "ideal" state scenarios when it comes to the operationalization of responsible AI principles, much of this will depend on organizational resources and maturity. The following principles, or some version of these principles, are thought to be important to consider and it is up to the organization to operationalize what is necessary to ensure alignment with these principles:

- Usefulness, Usability, & Efficacy
- Fairness & Bias Management
- Safety
- Privacy & Security
- Transparency
- Fairness/Bias

Operationalization may look like:

- Establishing requirements for risk-based validation and monitoring of usefulness, efficacy, safety, and bias/fairness (supported internally, by an external third-party, and/or vendor/developer).
- Standard documentation requirements for vendors/developers, standard disclosures/consent procedures for certain AI solutions, and use-case specific training and feedback requirements (transparency) (e.g. Applied Model Card and AI Intake Questionnaire).

Key Resources & Tools

- Responsible AI Lifecycle Management diagram in the CHAI Responsible AI Guide
- FDA AI Lifecycle
- ISO/IEC 5338:2023 Information Technology--Artificial Intelligence--AI System lifecycle processes
- CHAI Applied Model Card: Digital Format
  - CHAI Applied Model Card: Word Version with Instructions
  - CHAI's Applied Model Card is intended for an applied AI solution in context, not an AI model outside of an intended use/purpose (e.g. general foundation models, or AI platforms used to build AI solutions/applications.
- CHAI Draft Common Intake Questionnaire

Responsible Data Management & Use

Data management and use are important for all digital health technologies, but in particular for health AI solutions (See Subdomain 4.3: Responsible Data Management & Use Playbook). Even when data are de-identified, mis-use of data can pose the risk of re-identification and bias. Organizations should consider establishing the following to ensure privacy-preserving and responsible data management and use when it comes to AI solutions:

- Designate a data steward (individual or team) to review data and security compliance.
- Contractual controls around primary and secondary data use
- Vendor security requirements
- Structured data acquisition and provenance tracking
- Data quality and bias management
- De-identification practices
- Standardized data preparation procedures

Third Party Management

When working with third parties (research teams or external vendors), organizations should consider discussing contractual terms related to liability and transparency requirements, clearly define lifecycle roles and responsibilities (including training, bi-directional feedback, monitoring, and data use), and ensure oversight of vendor performance and risk (See Subdomain 4.4: Third Party Management Playbook).

Education, Training & Feedback

Safe and accountable AI necessitates appropriate training and feedback mechanisms. These may include:

- Workforce training (e.g. use case specific, general AI literacy)
- Protected and transparent feedback channels (between users, governance bodies, and developer/vendor when appropriate)
- Clear incident reporting responsibilities, stakeholder-specific transparency practices, patient privacy communication, and controlled testing and rollout processes.
- See Subdomain 4.5: Education, Training & Feedback Playbook

Baseline Operational Controls for Domain 2: Organizational Structures

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

Developing an AI Governance Committee and Assigning Critical Roles & Responsibilities

	Control 2.1: Ensure your organization has staff assigned to critical roles and responsibilities (as defined by your organization). These can be existing or new staff, and one staff member may hold multiple responsibilities (one-to-many).
	Control 2.2: Establish a formal committee structure bringing together all governance stakeholders

Effective AI governance depends on clearly defined human accountability and a structured forum for collective decision-making. It is therefore important to identify who will be responsible for each governance function, and establish a space for these responsible parties to engage in dedicated, cross-functional coordination. Governance committee structures may shift and evolve with changes in staffing, the organization's AI portfolio, or regulatory shifts, however adopting a structure that can scale with your organization, without adding unnecessary bottlenecks is valuable.

Implementation Guidance

Map Governance Functions to Roles

Ideally, organizations should adopt a formal governance structure that is responsible for risk-based, and organizationally appropriate oversight of health AI solutions involved at the very least, in:

- Direct or indirect patient care
- Care support services
- Care-relevant healthcare operations and administrative services

Governance structures should be comprised of multiple relevant stakeholders, with a designated individual who has appropriate technology expertise, ideally in AI.

The AI governance structure does not need to be its own standalone team, but should include individuals with the following expertise as appropriate:

- Executive leadership
- Finance
- Regulatory/Ethical compliance
- Information Technology
- Safety/Incident Reporting
- Relevant clinical/operational/administrative expertise
- Cybersecurity and Data Privacy Needs
- Stakeholder or liaison reflecting the needs of impacted populations (e.g. staff, providers, patients, caregivers, etc.)
- Finance

Note that with novel technologies, even individuals with these titles or expertise may not have a necessary grasp on AI-specific (predictive, generative, and agentic) considerations and knowledge. Consider getting a sense of where staff are at, and what kind of additional training or competency might be required in order to participate. See Subdomain 4.5: Education, Training, and Feedback Playbook for more on this topic.

Note and Examples

AI governance committees do not need to be newly formed committees. For example, one large multi-state health system chose to integrate their AI Governance into their Data Governance body. They added an AI Stewardship Role and created an AI Governance Council that works alongside the Data Governance Council.

Similarly, a large community health center, grounded their AI intake, review, and monitoring process into their Quality Improvement Team, where there were individuals with the skills required to effectively pilot, implement, and monitor new solutions.

At minimum, organizations must identify who is responsible for oversight, validation, monitoring, and data stewardship.

Action: Conduct an AI skills assessment of current clinical leadership and leadership in IT, Quality, Risk Management, and Clinical Informatics to identify gaps and needs. Identify who has the authority to introduce, approve, validate, turn off, and decommission AI solutions (either those developed internally or those procured externally). If gaps are identified, consider up-skilling existing staff or hiring additional staff as appropriate. To avoid conflict of interest, it is important that the validator of the solution is not also the developer of the solution.

Define the Committee Structure

The AIGC serves as the "institutional front door" for all AI initiatives, handling intake, risk-based oversight, and policy enforcement.

Action: Draft an AI governance committee charter, or update related committee charters to account for updates made to the process and AI relevant considerations.

Models of Governance Structure

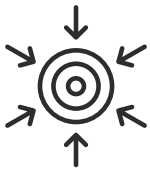

Centralized

A centralized AI governance structure concentrates decision-making authority, standards, and oversight within a single enterprise-level body or function (e.g., a central AI governance office or committee).

Core Characteristics

- One governing authority approves AI solutions and policies
- Standardized processes, templates, and controls
- Central ownership of risk assessments, monitoring, and reporting
- Limited local autonomy

Tradeoffs

- Slower decision-making
- Risk of bottlenecks
- Less sensitivity to local workflows

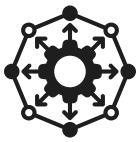

Decentralized

A decentralized AI governance structure places primary responsibility for AI oversight within individual departments, service lines, or business units, with minimal central coordination.

Core Characteristics

- Local teams select, deploy, and govern AI solutions
- Governance decisions tailored to specific workflows
- Minimal enterprise-wide standards
- High autonomy, low consistency

Tradeoffs

- Inconsistent standards and documentation
- Increased risk of “shadow AI”
- Harder to detect systemic risk or bias

Distributed

A distributed AI governance structure combines central standards and guardrails with local execution and accountability.

Core Characteristics

- Central body sets:
  - Policies
  - Risk tiers
  - Minimum requirements
- Local teams:
  - Own intended use
  - Manage day-to-day operations
  - Monitor local performance
- Shared templates, tools, and escalation pathways

Distributed governance models work best when teams share standardized artifacts, not just policies. These might include things like risk taxonomies, threat models, and evidence templates that allow for local autonomy without sacrificing consistency, auditability, or regulatory readiness.

Note: Governance Structures Evolve

There is no single governance model that fits every organization, and the right structure today may not be the right structure in two years. As organizations grow their AI portfolios, build internal expertise, and distribute AI adoption across departments, many find that their governance approach naturally shifts — often from a centralized model toward a more federated or hybrid one.

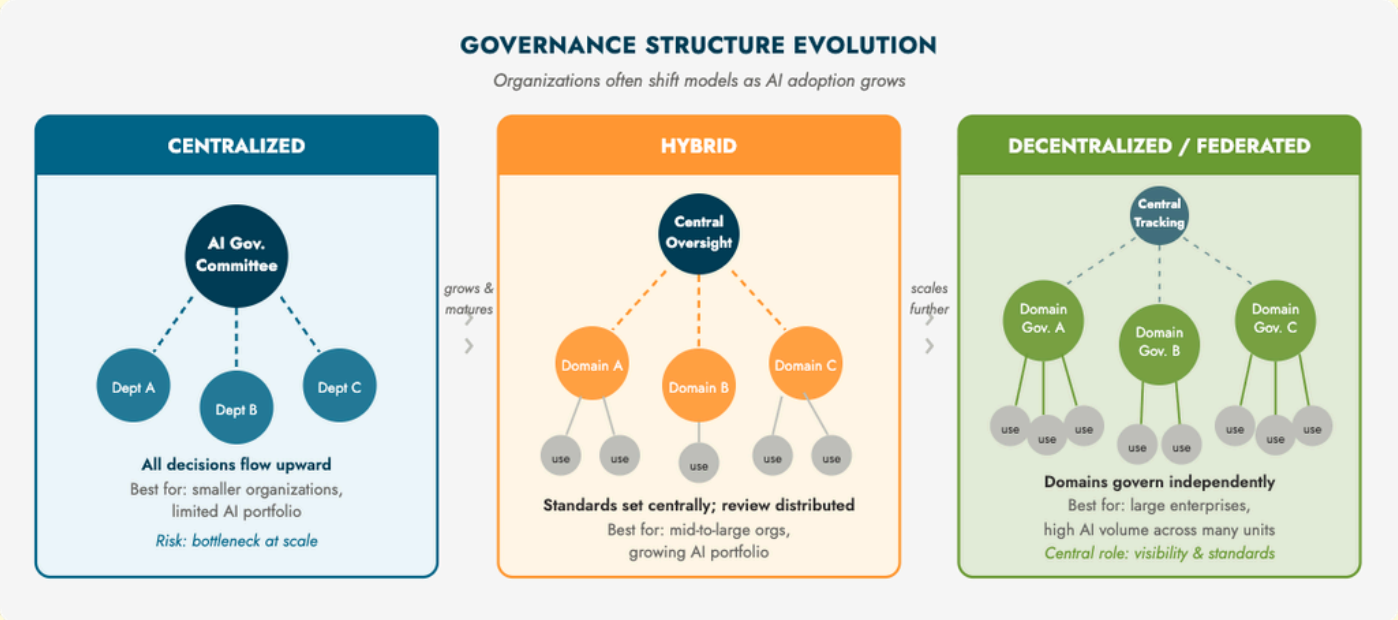

Organizations should periodically reassess their governance model as part of their standard policy review cycle, particularly following significant growth in AI adoption, major organizational changes, or shifts in regulatory expectations.

To read more about defining the scope of the governance structure, see the AI Policy Playbook.

Note: Considerations for Community Health Centers (CHCs) and Federally Qualified Health Centers (FQHCs)

Many smaller or less resourced health centers/systems may not have access to a team that covers all roles and responsibilities, or may have one individual that fulfills multiple roles. When using a risk-based approach, governance at it's core is meant to be an enabler by establishing faster approval paths for low-risk AI, and ensuring appropriate safeguards are in place for higher-risk AI.

Rather than forgo a governance structure or rely solely on vendor reports, Community Health Centers (CHCs) and Federally Qualified Health Centers (FQHCs) can approach governance in a distributed way across local, regional, and networked partners. They may also distribute governance roles across existing committees rather than standing up a new body.

Health Center Controlled Networks and Primary Care Associations can function as a type of distributed governance infrastructure, help provide stronger negotiation power during contracting, and maintain technical/training support. The following is a possible illustration of how roles and responsibilities may be distributed.

CHCs/FQHCs

- Identify relevant problems that can be addressed with AI solutions
- Make decisions on how clinical workflows are structured and modified with the introduction of new tools
- Retain local incident response processes
- Serve as final approval and enforcement of policies
- Local customization
- Can offer local clinical, patient, administrative, or operational expertise
- Decide whether and how to use each tool

PCAs

- Provide AI policy starter sets
- Provide risk assessment templates
- Establish vendor intake checklists
- Provide model card/procurement guidance
- Host office hours, trainings, and shared expert panels to improve network AI literacy, data science expertise, contractual assistance
- Share feedback and experience from multiple CHCs

HCCNs: Already manage shared infrastructure and data flows

- Participate in the shared evaluation and purchasing of commonly requested tools (along with PCAs where appropriate)
- Maintain shared AI inventories and track commonly deployed tools
- Flag known data/privacy risks or updates
- Assist in monitoring and de-identified incident aggregation through shared dashboards or logging tools
- Help develop role-based AI training and provide updates when tools change

This example model helps to reduce duplication and fragmentation across the network, increases leverage during negotiations, allows scarce expertise a wider reach, and provides flexibility for small and large organizations across the safety net. Most importantly, this means that governance quality and access to innovation are not limited by capacity.

Community Insight

A common challenge when setting up governance roles and accountability is that committee decisions often lack a clear execution path into engineering, security, and compliance workflows. Governance structures are most effective when paired with an operational layer that translates decisions into traceable tasks, controls, and evidence.

Key Tools & Resources

- AAMC: Artificial Intelligence Competencies Across the Learning Continuum (Coming Fall 2026; Targeted for Medical Education and Practice)
- International Association of Privacy Professionals: Artificial Intelligence Governance Professional Certification (Relevant for professionals tasked with developing AI governance and risk management in their operations--not healthcare specific)

Intake, Assessment, and Monitoring Framework

Control 2.3: Establish an intake, assessment, and monitoring process for thorough investigation of prospective solutions and mechanisms for incident detection, feedback, and follow-up.

Without a defined intake, assessment, and monitoring framework, governance becomes reactive — responding to failures rather than preventing them. The responsible adoption of AI is not a single decision, it is a continuous process. Before a solution is deployed, organizations need a structured way to evaluate it against clinical, operational, ethical, and technical criteria. After deployment, they need mechanisms to detect when something is going wrong, capture feedback from end users, and act on it before issues escalate. More information can be found in the Subdomain 4.1: Responsible Lifecycle Management & Use Playbook.

Implementation Guidance

Standard Intake Process

1. Require that internal developers and vendors complete a Model Card (or similar documentation) and Standard Intake Application
    - Consider using a tiered intake process, with a brief screener (or the model card alone) that includes all minimum threshold/non-negotiable items. This saves the vendor and healthcare delivery organization from needing to complete deeper intake questions or review.
    - **Model card:** Information in a standard model card (or similar documentation) can be used to streamline the initial information collected from vendors, creating a consistent process. It makes AI risks, responsibilities, and readiness visible early, improving safety, efficiency, and decision-making

with minimal added burden. The CHAI applied model card is intended for an applied AI solution in context, not an AI model outside of an intended use/purpose (e.g. general foundation models, or AI platforms used to build AI solutions/applications.

- Model cards alone are rarely enough for supporting downstream regulatory or risk workflows. Ensure that additional information requested aligns with what might be needed for proper risk categorization. This will likely include questions related to security, compliance, data, and clinical/operational risks. CHAI is in the process of developing a common intake guide and questionnaire (see Draft CHAI Intake Questionnaire here).

- Consider adding a required risk mitigation and monitoring plan in addition to the model card, especially for known risks specified on the model card.

Real-World Example

When it comes to tiered intake, a large academic medical center shared that they don't spend time on solutions that can't meet baseline requirements. For example, if their Enterprise AI Policy requires a minimum threshold for demographic representation in model training data, and a vendor can't share or meet it, they would want to surface that very early on and stop or redirect the developer or vendor before deeper review.

2. Categorize developer/vendor & AI Solution Risk (Pre-Deployment)
3. Identify Gaps in Vendor/developer information based on initial risk categorization exercise. Note that existing organizational security review processes can be updated to include an additional AI security risk evaluation.
4. Follow-up with vendors/developers by providing feedback and requesting clarification on gaps identified
5. Develop a plan for how the solution's effectiveness and workflow fit will be evaluated and monitored and who will be responsible. Note that there are many different levels of rigor when it comes to assessment and monitoring and that a risk-based approach should be used.
6. Develop a plan for how and under what conditions relevant parties will be informed about updates, adverse events, incidents, or feedback, as well as who is legally or contractually accountable (See Subdomain 4.4: Third Party Management and Subdomain 4.5: Education, Training, and Feedback Playbooks for more on these topics.) This should include a protocol for turning off or decommissioning the AI system.

Note: Looking Forward, Using AI to Help Govern AI

The governance structures described in this playbook are designed to be human-led, but that does not mean they need to be entirely human-executed. AI itself can play a meaningful role in making governance more efficient, consistent, and scalable, particularly as the volume of AI solutions under management grows. This is particularly true at institutions with large AI portfolios, where steps are already being taken to reduce bottlenecks and make governance processes efficient and effective.

As always a risk-based approach should be taken, and human oversight methods (in-the-loop, on-the-loop, or out-of-the-loop) should be determined accordingly. Organizations should consider how AI-powered tools can support the governance function across several areas:

Monitoring & Telemetry. Automated monitoring of deployed AI solutions can surface performance drift, unexpected outputs, or anomalous usage patterns, triggering alerts before issues escalate to the governance committee. This reduces the burden of manual post-deployment oversight on AI Owners and committee staff.

Intake & Pre-Review. AI can assist in screening new use case submissions, flagging gaps in documentation, completing model cards based on available information, checking proposed solutions against existing policy criteria, and triaging requests by risk level before human reviewers engage. This accelerates the intake process and focuses reviewer time on higher-complexity decisions.

Policy Alignment Checks. As policies evolve, AI solutions can help assess whether active solutions remain aligned with updated requirements, identifying use cases that may warrant re-review without requiring staff to manually audit every entry in the AI inventory.

Audit & Reporting. Automated reporting on governance activity, such as: submission volumes, review timelines, approval rates, outstanding action items, reduces administrative overhead and provides leadership with real-time visibility into the governance pipeline.

As these capabilities mature, organizations that invest in AI-enabled governance infrastructure will be better positioned to scale oversight without proportionally scaling staff. The governance committee's role shifts from execution to exception handling, focusing human judgment where it matters most.

Communication and Resolution Protocols

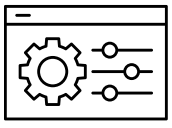

Control 2.4: Establish protocols for escalating and communicating concerns/incidents for resolution to the governance structure

This section specifically refers to feedback and incident reporting from the individual(s) responsible for monitoring the AI solution, and the individual responsible for managing and approving the use of the AI solution. The process will likely differ depending on whether your organization has a centralized, decentralized or distributed AI governance structure. (See Subdomain 4.5: Education, Training, and Feedback for user, vendor, and organization feedback and incident reporting guidance.)

While most issues should be communicated to the AI system owner, there may be instances when the governance structure may need to be informed about the resolution of a concern or incident, or when escalation may be needed.

Implementation Guidance

Task 1: Define What Triggers Escalation to the Governance Structure

Not every performance issue or user complaint warrants governance-level attention. Defining escalation triggers in advance prevents both under-reporting (monitors unsure what rises to that level) and over-reporting (every minor anomaly flooding the committee). Essentially, answer the question: Under what conditions does the governing body need to know when something goes wrong and assist in its root-cause-analysis or resolution? Organizations should establish at minimum two or three escalation tiers. See below for an example.

Note: Escalation Tiers

A *routine concern*: minor output anomaly, isolated user complaint, slight performance drift — is handled and logged at the operational level by the AI owner or clinical champion without immediate governance notification.

A *moderate concern*: unexplained pattern of performance degradation, repeated user complaints pointing to a systematic issue, evidence the tool is being used outside its intended scope — is communicated to the designated governance contact (often the committee chair or a standing governance liaison) within a defined window (e.g., 48-72 hours), with a decision on whether to place it on the next meeting agenda.

A *serious incident:* patient safety event involving AI output, suspected differential harm across patient subgroups, security or data breach involving the AI system, or AI behaving in ways that appear to deviate substantially from its intended function — requires immediate escalation to the governance structure plus any other required functions (compliance, legal, CISO, patient safety). Organizations should define "immediate" as a specific timeframe, not just a word.

Task 2: Establish a Designated Escalation Recipient and Pathway

"Escalate to the governance committee" is not actionable without a named contact. Identify who receives escalations on behalf of the governance structure — this is typically the committee chair, a clinical informatics lead, or a designated AI governance officer. That person is responsible for acknowledging receipt, making an initial triage determination, and deciding whether the matter requires convening the full committee, delegating to a subgroup, or monitoring with a follow-up timeline.

Define the mechanism: escalation should follow a consistent channel — a secure form submission, a structured email to a governance inbox, or a report in the AI inventory system. Consistency matters both for governance and for documentation.

Task 3: Define a Minimum Standard Escalation Report

The person escalating should not have to determine what information the governance body needs. Provide a template or required fields for any escalation communication (e.g. like a Patient Safety & Incident Reporting

System). See for additional guidance.

Task 4: Define What the Governance Structure Does Upon Receipt

Define the governance body's obligations upon receiving an escalation: a required acknowledgment to the escalating party within a defined timeframe; a decision pathway with defined options (add to next agenda, convene emergency session, delegate to a subgroup, suspend the AI tool pending investigation); and documentation of the disposition. The authority to suspend or restrict an AI tool pending review should be explicitly named — who has that authority, and under what circumstances can they exercise it without waiting for a full committee meeting.

Task 5: Document, Close the Loop, and Learn

Organizations should consider logging in a dedicated incident register (which can be a tab in the AI inventory). The escalating party should receive a disposition notification — what the governance body decided and why. Patterns across incidents should be reviewed at least annually to inform whether risk tiers for specific AI solutions need to be revised, whether monitoring responsibilities are appropriately staffed, and whether the escalation protocol itself needs adjustment.

Case Studies: Example Governance Structures

Explore these real-world examples of how organizations are successfully approaching health AI governance. These case studies offer valuable insights and practical strategies that you can adapt to your own organization.

Real-World Example #1:

A 13-hospital system that includes community hospitals and academic tertiary medical centers shared that AI adoption is decentralized within their organization. Novel AI solutions are brought forth by individual business units as they identify their own needs, and are approached on a case by case basis. In their context, solutions are evaluated and approved by the Information Support group and governance is overseen by the Chief Information Officer & Chief Medical Information Officer, providing both technology and clinical perspectives. Formal, standard governance processes are still evolving at their organization.

Real-World Example #2: Johns Hopkins

Approval process and governance structure steps when a vendor approaches:

1. Identify an internal sponsor (clinical or operational stakeholder): prioritized
2. Rigorous intake process: evaluation of functionality, impact, ethical considerations, legal and compliance risks, usability concerns, privacy & security risks, and return on investment. Required documentation from vendors include model cards and data cards.
3. Route to one of 3 subcommittees for clinical, imaging, or operational review.
4. Implement success metrics based on clinical or operational goals

Source: Johns Hopkins: An AI Success Story

Example #3: Mayo Clinic

Approval process and governance structure steps when a vendor/innovator approaches:

1. Established an independent multidisciplinary physician-led board to assess risk and regulatory applicability for digital health technologies (DHTs), including AI.
2. Innovators submit their solutions through a centralized intake process.
3. Tools that clearly do not meet the definition of a medical device (administrative support or low-risk tools) go through a fast-track analysis (few hours of team prep + 30 min review, with decision finalized offline by the board).
4. Tools that may be subject to FDA regulations go through a two week detailed analysis, which includes iterative communication between the product team and regulatory experts.
5. For tools that fall into uncertain/unclear "gray areas", consensus is reached by using weighing factors such as time-criticality (Is the tool used in high-stress scenarios where output cannot be easily verified?), explainability (Is the output transparent enough for independent verification?), and automation bias (Is the tool likely to trigger over-reliance by the user?)
6. Issue formal report with regulatory classification, risk controls aligned with international standards (ISO 13485 & ISO 14971), and lifecycle requirements for ongoing monitoring and traceability.

Source: https://www.sciencedirect.com/science/article/pii/S294976122400083X

Example #4: TrueCare

This example reflects how TrueCare, a nonprofit community health center, approaches AI governance, with a focus on organizational structure:

(1) Board Alignment & Education — Establish strategic AI priorities with the Board and maintain a regular cadence of updates and education to support informed oversight and decision-making.

(2) Standard Intake & Triage — Route all AI requests (or use cases) through a single intake process. Classify as clinical or administrative, document the business need, and assign a business owner before any evaluation begins.

(3) Clinical Pathway (EHR governance) — Direct clinical use cases through EHR governance for structured review, including patient safety, clinical validation, and equity assessment.

(4) Administrative Pathway (leadership review) — Route operational use cases through senior leadership team to validate the business case, alignment, and readiness.

(5) Chiefs Approval (required for all AI use) — All AI use cases and tools—clinical and administrative—require final review and approval by the Chiefs prior to pilot, procurement, or deployment.

(6) KPIs & Monitoring — Define success metrics upfront (e.g., adoption, efficiency, quality). Require ongoing reporting to assess performance and inform decisions to scale, adjust, or sunset.

(7) Policies & Practical Education — Maintain clear, concise AI policies that are easy to understand and apply. Pair with required, annual, role-based education to ensure teams feel confident—not intimidated—in using AI safely, ethically and responsibly.

Source: Chief Innovation Officer, TrueCare

Related Challenges

Below are challenges that organizations identified in by HDOs when working to implement organizational resource controls. These are real barriers, and recognizing them is the first step to addressing them.

Challenge/Gap	Details
AI Readiness Assessment	Difficulty applying assessments without policies or structure; lack of standards. Lack of standardized tools for readiness and assessments.
Staffing and Expertise	Lack of qualified staff, especially in smaller orgs and clinical informatics.
AI Governance Committee	Challenges in stakeholder alignment, avoiding fatigue, and ensuring actionable governance.
Resource Constraints/Staffing	Difficult to dedicate staff or justify new non-revenue generating roles.
Reporting Channels and Culture	Ensuring secure, anonymous, and trusted reporting systems.
Non-Retaliation Policies	Need to integrate with existing safety protocols; AI concerns can create unique vulnerabilities.
Escalation Pathways and Root Cause Analysis	Lack of AI-specific root cause analysis tools and timely escalation mechanisms.

Tools & Resources

CHAI Workbook to help design your organizational governance structure |This workbook is designed for organizations of different sizes to guide the development of their AI governance structures

CHAI Role Description Guide |A non-exhausting set of roles and responsibilities for AI governance

Additional Resources (Non-Exhaustive)

- DiME Implementation RACI Matrix |Core personnel to include on an implementation team at HDOs
- DiME Health AI Readiness Assessment |Self-assessment around governance, infrastructure, and capabilities for HDOs
- Gartner AI Maturity Model and AI Roadmap Toolkit | Gartner | CIO Tool for governance maturity assessment and strategy development (sector agnostic)

- The MITRE AI Maturity Model and Organizational Assessment Tool Guide | MITRE | Sector Agnostic AI Maturity Model
- DiME Health AI Evaluation Guide |High level basic questions that can be useful during a tier 1 intake for HDOs

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