Domain 3: Organizational Resources 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

AI systems do not operate in a vacuum. They require the right data infrastructure, tools, computing capacity, human expertise, and documentation practices to function safely and responsibly. Without deliberate resource governance, organizations risk deploying AI on unstable foundations—exposing patients, staff, and institutions to avoidable harm.

The **Domain 3: Organizational Resources** playbook is about ensuring your organization has intentionally accounted for the resources AI systems depend on, and that those resources are governed, documented, and secured.

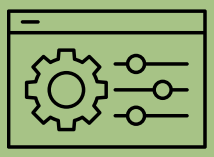

Baseline Operational Controls for Domain 3: Organizational Resources

Category	Control
Data Resources	Identify and document security infrastructure and methods (HITRUST, SOC 2, NIST CSF, CISA) available and needed to protect data; including methods for encrypting data at rest and in transit.
Data Resources	Identify and document the tooling and data resources available and needed to conduct regular security assessments and vulnerability scans to identify and address risks.
Resource Documentation	Maintain comprehensive records of AI-related assets including datasets, models, and tools to ensure traceability.
Tooling Resources	Identify and document tools, resources, and methods available and needed to implement AI cybersecurity evaluation (penetration testing, adversarial robustness, etc.).
Tooling Resources	Identify and document tools and resources available and needed for standard documentation of all AI systems and for integrating them into an AI inventory, such as adopting a Model Card.
Tooling Resources	Identify and document the tools and resources available and needed to implement and maintain an AI inventory registry to track internal and third party solutions.
Tooling Resources	Identify and document the tools, resources, and infrastructure available and needed to carry out validation testing and monitoring for relevant AI models, including responsible AI baseline tests.
Computing Resources	Identify and document the computing resources available and needed to conduct recurring cybersecurity audits of AI infrastructure; enforce access controls, and HIPAA compliance.
*Human Resources	Establish an AI workforce competency framework (or related document) that defines: roles and responsibilities of AI governance personnel, minimum qualifications and/or certification. Include ongoing training programs to update personnel on governance concepts such as AI safety and bias mitigation.
*Human Resources	Identify when clinicians and staff need to be trained on specific AI solutions prior to implementation; also, determine if certain AI solutions require regular training/competency assessments.

* controls related to Human resources did not originally meet cutoff (>=67%) for inclusion based on “must-have” determination from both governance platforms and health systems, however during playbook work group feedback (Phase 2 feedback), it was determined that these should be included. These are therefore highly recommended but not required.

Risk-Based Control Application

Note that when identifying and documenting tools, resources, and processes below, organizations should take a risk-based approach (as recommended for all controls across all playbooks). Not all solutions will need the same level of rigor in documentation, evaluation methods, monitoring, etc. Taking a risk-based approach reduces burden on governance teams and ensures that resources are appropriately allocated. See CHAI's Risk Tool here.

Data Resources

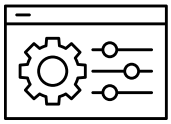

Control 3.1: Identify and document security infrastructure and methods (e.g. HITRUST, SOC 2, NIST CSF, CISA) available and needed to protect data; including methods for encrypting data at rest and in transit.

Healthcare data is among the most sensitive information in existence. AI systems routinely access, process, and act on patient data—making robust security infrastructure a non-negotiable foundation for responsible AI governance. Beyond HIPAA requirements, healthcare organizations must align their AI security posture with established cybersecurity frameworks that address modern threats, including AI-specific risks such as data poisoning, model inversion attacks, and inadvertent re-identification.

Implementation Guidance

Task 1: Establish a Security Baseline Aligned with Recognized Frameworks

Identify and document the security infrastructure your organization currently uses for AI-related data: encryption solutions, access control systems, network segmentation, cloud security configurations, and any certifications or attestations. Map this infrastructure against at least one recognized framework. Even if your organization is not yet certified, these frameworks provide concrete implementation checklists you can work toward progressively. Start with a gap analysis against HIPAA Security Rule technical safeguards as a minimum baseline and obtaining and reviewing vendor's security posture and certifications.

The most commonly used frameworks include: HHS 405b HICP, HITRUST CSF, NIST Cybersecurity Framework (CSF) 2.0, SOC 2 Type II (third-party audited vendor assurance), HIPPA Security Risk Assessment, and CISA guidance (AI-specific threat vectors). Document your current state and note gaps.

Action: Complete a security gap analysis comparing your current AI data handling against HIPAA Security Rule requirements and one additional framework or approach. Document your current state and prioritize gaps by risk level.

Task 2: Document Encryption Standards for AI Data Flows

For each AI system, document:

(1) What data is stored or processed (including PHI sensitivity level)

(2) Encryption method at rest (AES-256 or equivalent)

(3) Encryption in transit (TLS 1.2 or 1.3 for all API connections)

(4) Whether encryption extends to model artifacts, training data backups, and inference logs. For vendor-managed AI, obtain and document vendor encryption documentation as part of procurement. Flag any data flows where encryption standards are not yet confirmed.

Task 3: Address AI-Specific Data Security Risks

AI systems introduce security risks not captured by traditional IT security frameworks. These include:

(1) Re-identification risk: de-identified data used in AI training can sometimes be re-identified via model outputs or membership inference attacks.

(2) Data poisoning: adversarial actors may manipulate training data to degrade model performance or introduce bias.

(3) Model inversion attacks: adversaries may extract sensitive training data through repeated model queries.

(4) Indirect Prompt Injection: Adversarial content planted in clinical text is a form of prompt injection. It may include malicious or accidental instructions embedded in free-text fields (chief complaints, patient messages, referral notes, prior visit summaries) that are later processed by an AI system. The AI may treat the embedded instructions as legitimate guidance and act on them, bypassing intended behavior.

> **Agentic AI raises the stakes:** As AI moves from passive summarization toward taking actions (ordering, routing, documenting), the consequences of successful injection escalate from misleading outputs to unauthorized clinical actions.

Note: Community Health Centers (CHCs)

Cybersecurity certification often requires meaningful upfront investment and internal documentation capacity that many CHCs simply don't have. SOC 2 Type II is largely a vendor credential, not a healthcare delivery credential. Formal NIST CSF adoption at a documented, structured level requires IT staff time and governance capacity that some small CHCs lack.

What CHCs actually rely on

The most common security posture for FQHCs and smaller CHCs include:

HRSA's annual HIPAA Security Risk Assessment (SRA) is legally required as a condition of federal funding, so it functions as the de facto baseline security framework for most CHCs. This is often their most rigorous annual security touchpoint.

EHR vendor cloud security carries a substantial portion of the security burden — most CHCs use hosted or cloud-based EHR environments (athenahealth, eClinicalWorks, Epic, Greenway), and that vendor's security posture (including SOC 2, HITRUST, encryption standards) largely defines the CHC's security posture for PHI. This is a compensating control, not a substitute for organizational security governance, but it's the practical reality.

HHS 405(d) Health Industry Cybersecurity Practices (HICP) is arguably the most appropriate and realistic framework for smaller healthcare organizations. It was explicitly designed with limited-resource providers in mind, provides a tiered approach, and is free. CHCs that do have some IT security maturity tend to reference HICP rather than HITRUST or NIST CSF.

Health Center Controlled Networks (HCCNs) and Primary Care Associations (PCAs) in some states provide shared IT security services, security risk assessment tools, educational materials, and shared purchasing of managed security services. This is one of the more effective models for smaller CHCs achieving meaningful security infrastructure without bearing the full cost individually.

Managed Security Service Providers (MSSPs) are increasingly common among mid-size CHCs, particularly for functions like vulnerability scanning, endpoint detection, and security monitoring. This outsources the capability rather than building it internally.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A 25-bed critical access hospital leverages its EHR vendor's built-in HIPAA-compliant cloud environment as its primary AI security layer, supplementing with annual HIPAA Security Rule self-assessments and free NIST CSF resources. The hospital designates its IT director as the AI security lead and documents a simple data flow map for its active AI solutions. The organization may also partner with their HCCN to or MSSP for help here.	A 400-bed regional system implements HITRUST e1 Certification (an entry-level HITRUST program designed for organizations with limited resources) as a stepping stone toward full HITRUST CSF Certification. It deploys an encryption management solution across all AI data pipelines and conducts quarterly reviews of its AI security posture alongside existing IT governance activities.	An academic medical center with 30+ AI solutions pursues full HITRUST CSF Certification, engages a third-party firm for annual SOC 2 Type II audits covering AI infrastructure, and maintains a dedicated AI cybersecurity function within its Information Security Office. It has implemented automated monitoring for anomalous AI data access patterns and maintains an AI security incident response plan aligned with its broader CISO program.

Key Resources & Tools

NIST Cybersecurity Framework (CSF) 2.0 — Free | Comprehensive framework with healthcare implementation guides

HITRUST CSF — Healthcare-Specific Framework | Integrates HIPAA, NIST, ISO, and more

CISA Guidelines | AI-specific security guidance

OWASP AI Security and Privacy Guide | Open-source AI threat library

HHS Office for Civil Rights — HIPAA Security Rule Guidance | Regulatory baseline

Microsoft Rural Health Resiliency Program | Free AI and cloud readiness assessments for rural hospitals and CAHs; supports NIST AI RMF and governance committee setup

HHS 405b HICP | The HHS 405(d) Health Industry Cybersecurity Practices (HICP) is a free, federally developed resource that provides practical, tiered cybersecurity guidance specifically designed for healthcare organizations of all sizes, including those with limited IT resources.

NACHC AI Action Guide | Practical 7-page guide for community health centers on AI implementation and governance (National Association of Community Health Centers)

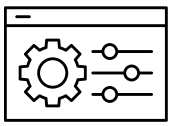

Control 3.2: Identify and document the tooling and data resources available and needed to conduct regular security assessments and vulnerability scans to identify and address risks.

Deploying AI in clinical environments creates a dynamic attack surface that must be continuously evaluated. Before an organization can conduct effective security assessments, it must understand what tools and data access are available—and identify gaps that need to be addressed. This control requires organizations to take stock of their current scanning and assessment capabilities, document those resources, and plan for gaps.

Implementation Guidance

Task 1: Inventory Existing Security Assessment Tools and Data Access

Document what security assessment and vulnerability scanning capabilities your organization currently has: commercial tools (e.g., Tenable, Qualys), open-source options (e.g. OpenVAS, NIST NVD), access to third-party testing firms, and which teams can perform assessments. Also document which AI-connected data environments and systems can be scanned. Note any access restrictions—particularly vendor-managed systems where scanning rights may be limited.

Action: Complete a capabilities inventory that lists available tools, their licensing status, scope of coverage, and responsible team. Identify AI systems or infrastructure that cannot currently be scanned and flag these as gaps.

Task 2: Identify Gaps and Resource Needs

Compare your current scanning capabilities against your AI system inventory (see Control 6). For each AI system, determine: Is it scanned? If not, why not—and what would be needed to include it? Common gaps include: (1) No scanning rights for vendor-managed cloud AI (address through contract language at procurement), (2) No staff expertise for AI-specific vulnerability assessment, (3) No budget for third-party penetration testing. Document these gaps and prioritize remediation by risk level.

Action: Produce a gap assessment document that maps each AI system or infrastructure component to available scanning coverage, and identifies what additional tools, access, or expertise would be needed to close gaps.

Task 3: Establish and Document an Assessment Cadence

Define and document the minimum assessment cadence for AI systems: prior to deployment, annually for all active AI systems, following significant changes, and after any security incident. For vulnerability scans, many organizations run automated scans monthly or quarterly. Where assessments cannot yet be conducted directly, document compensating controls (e.g., obtaining vendor SOC 2 reports, requiring third-party pen test summaries).

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Approaches to security assessments vary significantly based on internal expertise and resources.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small hospital documents that it relies on vendor-provided SOC 2 reports for all AI solutions (no internal penetration testing capability). It uses OWASP free resources to brief IT staff on AI-specific threats and has added a contract clause for future AI vendors requiring security incident notification within 48 hours. It documents that adversarial robustness testing is a gap for its two clinical AI solutions and records this as a medium-term capability to build.	A regional health system documents an annual third-party penetration test covering its three highest-risk AI systems. It maintains an AI Security Incident Response Playbook and records which AI solutions are covered by the annual test scope versus those relying on vendor SOC 2 attestation. A gap assessment identifies that one newer NLP-based AI tool for clinical documentation has not yet been included in penetration testing scope—added to next year's testing plan.	A large academic medical center documents a dedicated AI security function with continuous API security monitoring, quarterly adversarial robustness assessments for high-risk AI, and participation in sector-wide AI security threat intelligence sharing. It maintains a formal AI Security Policy. A gap assessment for a newly acquired AI-powered diagnostic tool identifies that vendor-side model theft protections have not been documented—a 30-day remediation item assigned to the CISO.

Key Tools & Resources

NIST National Vulnerability Database (NVD) | Free vulnerability reference database

CISA Known Exploited Vulnerabilities Catalog | Free; prioritized vulnerability list

HITRUST Third-Party Risk Management Program | Vendor assurance framework

HHS 405(d) Health Industry Cybersecurity Practices | Free healthcare-specific cybersecurity guidance

Resource Documentation & Tooling Resources

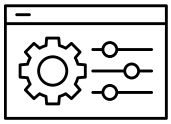

Control 3.3: Maintain comprehensive records of AI-related assets including datasets, models, and tools to ensure traceability.

Resource documentation is the connective tissue of AI governance. Model cards document individual AI systems. The inventory registry tracks what is deployed. Security assessments document vulnerabilities and remediations. But governance also requires documentation of the underlying assets these systems depend on: the datasets used for training and validation, the tools used for development, evaluation, and monitoring, and the versions and configurations that determined system behavior at any point in time.

Without comprehensive asset records, organizations cannot: trace the root cause of an AI-related adverse event, demonstrate compliance to regulators, conduct meaningful audits, or make informed decisions about AI system updates or decommissioning.

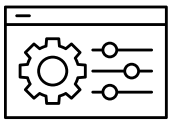

Control 3.4: Identify and document tools, resources, and methods available and needed to implement AI cybersecurity evaluation (penetration testing, adversarial robustness, etc.).

AI systems present a distinct and expanding cybersecurity attack surface that extends beyond traditional IT security. Adversaries can attempt to manipulate model outputs (adversarial attacks), extract training data (model inversion), steal model intellectual property (model theft), or cause clinical AI to malfunction at critical moments. Before an organization can conduct AI cybersecurity evaluation, it must know what tools and methods are available—and identify where it has gaps. Tools for AI cybersecurity evaluation are still maturing, making intentional resource documentation even more important.

Implementation Guidance

Task 1: Document Available AI Cybersecurity Evaluation Tools and Expertise

Identify and document what tools and expertise are currently available to conduct AI cybersecurity evaluations: penetration testing capabilities (internal team, external firm), adversarial robustness testing frameworks, API security testing tools, and relevant staff expertise. Note which AI systems can currently be tested and which cannot (e.g., vendor-managed systems where testing rights have not been negotiated).

Action: Produce a cybersecurity evaluation capabilities inventory, noting: available tools, who can use them, which AI systems are in scope, and testing rights status for each vendor.

Task 2: Define Required Evaluation Scope and Identify Gaps

Using a risk-based approach, determine what cybersecurity evaluation should cover for each AI system: (1) Penetration testing of inference endpoints and APIs, (2) Adversarial robustness (particularly relevant for image-based AI such as radiology), (3) API security—authentication, rate limiting, logging, (4) Data encryption verification throughout the AI pipeline, (5) Model theft protections, (6) Input validation against malformed or out-of-distribution inputs. Compare this required scope against available tools and capabilities, and document gaps. Prioritize gaps for high-risk clinical AI systems.

Task 3: Document Vendor Evaluation Arrangements and Contract Requirements

For vendor-managed AI systems, document what cybersecurity evaluation evidence the vendor currently provides (SOC 2 reports, pen test summaries, vulnerability disclosures) and whether contractual testing rights have been negotiated. Where gaps exist—vendors that restrict testing and provide no assessment evidence—document these as procurement remediation items. Include AI security evaluation rights as a required contract provision going forward.

Key Tools & Resources

OWASP Machine Learning Security Top 10 | Open-source ML threat taxonomy

MITRE ATLAS — Adversarial Threat Landscape for AI Systems | AI-specific threat matrix (analogous to MITRE ATT&CK)

CISA — Securing AI: Similar Approaches, Familiar Challenges | Federal AI security guidance

NIST AI RMF Playbook — GOVERN and MEASURE | AI risk management including cybersecurity dimensions

CHAI Privacy & Cybersecurity Profile based on NIST | Page 130 of CHAI's Responsible AI Guide

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Approaches to security assessments vary significantly based on internal expertise and resources.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small hospital or FQHC uses a single SharePoint site as its AI documentation hub, with folders for each AI system containing: the model card, the vendor contract, the security assessment summary, and a brief data source description. The documentation owner is the clinical informatics or IT lead. The organization uses its existing document retention policy as a starting point, with compliance officer review to confirm AI records applicability.	A regional health system implements a structured documentation template for each AI system's asset record, covering datasets, tools, and versions. Templates are stored in SharePoint and linked from the AI inventory registry. A semi-annual documentation audit is conducted by the clinical informatics team, with gaps reported to the AI Governance Committee.	An academic medical center implements a formal AI Asset Management System integrating model cards, dataset documentation, version histories, and audit logs. Automated version control captures changes to model artifacts and configurations.

Key Resources & Tools

Datasheets for Datasets (Gebru et al., 2021) — Communications of the ACM | Foundational framework for dataset documentation

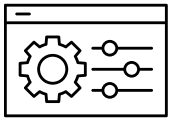

Control 3.5: Identify and document tools and resources available and needed for standard documentation of all AI systems and for integrating them into an AI inventory, such as adopting a Model Card.

A Model Card is a structured documentation artifact that captures critical information about an AI system: what it does, how it performs, what data it was trained on, how it should (and should not) be used, and what its known limitations are. In healthcare, model cards are an essential tool for enabling clinicians, governance committees, and compliance teams to make informed decisions about AI use.

CHAI has developed an Applied Model Card (AMC) specifically designed for HDOs, providing a standardized, healthcare-contextualized template that reduces the documentation burden while ensuring coverage of key governance fields.

Note: Adoption Variability

Workshop feedback indicated that many health systems want to adopt model cards but face practical barriers: vendor resistance to sharing information, lack of standardized formats, and capacity constraints. Governance platforms overwhelmingly see model cards as essential infrastructure for ecosystem-wide accountability.

Implementation Guidance

Task 1: Adopt and Customize a Standard Model Card Template

Start with an existing model card template rather than building from scratch. CHAI's Applied Model Card (AMC) is recommended as it is designed specifically for healthcare delivery contexts. At minimum, each model card should capture: Model ID and version, Intended use and primary end users, Out-of-scope uses, Training data sources and any known biases, Performance metrics (overall and across relevant subgroups), Safety considerations and known failure modes, Monitoring and alerting information, and Vendor/developer contact information.

Task 2: Establish a Model Card Completion Workflow

Designate model card reviewers (staff responsible for reviewing and updating cards for each AI system). Initial model cards should be submitted by the developer or vendor, while intial review may be the responsibility of the AI governance lead or equivalent, and updates may be the responsibility of a health informatics analyst or whoever is closest to the monitoring of the AI solution. For vendor-supplied AI, require model card information as part of procurement—many vendors will provide this in response to formal requests even if they do not proactively share it. For internally developed AI, require card completion as a mandatory step in your AI development or approval workflow. Consider making model card completion status a prerequisite for deployment approval (i.e., no model card = no deployment clearance).

See the CHAI Model Card Registry (Beta), for examples of completed model cards. You can also submit model cards or request that vendors submit model cards to the registry here.

Task 3: Integrate Model Cards with Your AI Inventory

Model cards are most useful when linked to your AI inventory registry (see Control 6) rather than maintained as standalone documents. This integration allows governance staff to quickly access documentation during vendor reviews, incident investigations, or AI governance committee meetings. Even a simple shared drive organized by model name and linked from your inventory spreadsheet is a meaningful first step. This integration can also make it easier to update documentation as new versions are deployed.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small hospital starts with a Model Card template and requires its two AI vendors to complete key fields as part of contract renewal. Cards are stored in a shared SharePoint folder accessible to the IT director, CMO, and compliance officer. The hospital uses a simple checklist to verify completion before approving any new AI tool.	A regional health system builds a model card completion requirement into its AI procurement checklist and assigns a clinical informatics analyst as model card coordinator. It uses a standardized Microsoft Forms-based intake that populates a SharePoint list (serving as a lightweight AI inventory), with model cards stored as linked PDFs. Incomplete cards trigger an automated reminder to the requesting department.	An academic medical center integrates a standard model card's fields into its MLOps platform (e.g., Azure ML, AWS SageMaker, or Domino Data Lab), enabling automated population of select fields from model training metadata. Designated model card reviewers in each clinical department review and attest to cards before deployment. The institution has submitted select models to the CHAI National Registry, establishing a public accountability record.

Examples

Michigan Medicine — Academic Medical Center

Michigan Medicine's Clinical Integration Center (CIC) has published a Model Card Library documenting AI models approved for use across the health system. The library, hosted on an internal SharePoint list, includes 24 models and captures standardized information about each AI system's purpose, performance metrics, ethical considerations, and limitations. The CIC requires model card completion as part of the endorsement process for any AI system brought forward for clinical use.

Source: Health AI Partnership, "Building Transparency: Artificial Intelligence Model Cards Inventory" (May 29, 2025) — healthaipartnership.org

Mercy - Large Multi-State Health System

Mercy's Responsible AI team stated that their procurement process requires completed model cards from vendors and internal developers of AI solutions procured or used within their system.

Key Tools & Resources

CHAI Applied Model Card (AMC) | Healthcare-specific model card template developed and maintained by CHAI (Digital Model Card Submission Link)

CHAI Model Card Registry (Beta)

Google Model Cards (Original Framework) | Foundation documentation for model card concept

FDA Transparency for Machine Learning-Enabled Medical Devices | Regulatory guidance on AI documentation

Health AI Partnership — AI Model Cards Inventory | Case study and guidance on building a model card library (features Michigan Medicine CIC)

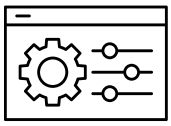

Control 3.6: Identify and document the tools and resources available and needed to implement and maintain an AI inventory registry to track internal and third party solutions.

The AI inventory registry is the foundational operational tool for AI governance, but before it can be built and maintained, organizations must identify what tools and infrastructure are available to support it. This control requires documenting current inventory capabilities, identifying gaps, and defining what resources would be needed to achieve a sustainable, current, and comprehensive registry. Organizations flagged shadow AI, unauthorized or undisclosed use of AI solutions, as a significant governance challenge that inventory processes must address. Some organizations choose not to deploy clinical or operational AI solutions, until they are fully registered in a centralized AI inventory.

Implementation Guidance

Task 1: Define the Scope and Fields of Your Inventory

Document what your AI inventory should include. At minimum, capture: Model/system name and unique identifier, Vendor or developer name and contact, Intended use and clinical/operational domain, Deployment status, Primary end users, Data inputs (including PHI sensitivity), Risk tier, Contract/license renewal date, Link to model card or documentation, Governance approval status and date, and Last review date. Document whether your inventory covers only clinical AI or also operational/administrative AI—many organizations start broad and refine scope over time.

Task 2: Identify and Document Available Inventory Tools

Document what tools are available to support your inventory: a shared spreadsheet (Excel or Google Sheets), SharePoint-based list, purpose-built AI governance platform (see Platforms in CHAI's Partnership Ecosystem), or integration with an existing CMDB or vendor management system. Document the current tool in use, its limitations, and whether it meets the key capabilities needed: searchable and accessible to governance stakeholders, version-controlled, linked to contract management for renewal alerts, and integrated with model cards. A well-maintained spreadsheet is a legitimate starting point.

Task 3: Document Shadow AI Detection Mechanisms and Intake Process

An AI inventory should include a documented process for identifying and auditing shadow AI use. Mechanisms can include: (1) A required intake process--no AI solution may be used without being registered, (2) A regular audit cycle (at minimum annually) to verify entries are current and detect undisclosed use, (3) A mechanism for staff to report AI solutions they are using that may not be in the inventory (tie to AI policy dissemination and shadow AI awareness efforts). Document who owns the intake process, how it is communicated to staff, and how compliance is monitored. Include inventory review as a standing agenda item for your AI Governance Committee.

Real-World Example

Intermountain Health — Large Integrated Health System

Intermountain Health manages a centralized AI inventory of more than 300 active AI solutions and over 4,000 automated processes. The organization established a Data Science and Artificial Intelligence Center of Excellence to oversee AI governance, maintain inventory currency, and ensure ethical standards are applied consistently across deployments. Intermountain's governance-led approach is explicitly designed to scale AI oversight across a large, geographically distributed system.

Source: Healthcare IT News (March 2025); Becker's Hospital Review, "Intermountain turns AI into a tailwind" (November 2025) — beckershospitalreview.com

AltaMed Health Services — Federally Qualified Health Center (FQHC)

AltaMed Health Services, a large FQHC network serving Southern California, integrated an AI-based predictive algorithm directly within its EHR workflow. The initiative is led by AltaMed's Medical Director of Clinical Informatics, who oversees the process for vetting and approving new AI solutions, a role analogous to an AI inventory governance function. AltaMed's experience illustrates how FQHCs with dedicated clinical informatics leadership can implement structured AI review processes at scale.

Source: Healthcare IT News, "One FQHC's AI story - from product research to results" (November 18, 2025) — healthcareitnews.com

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A critical access hospital or FQHC documents that it maintains a shared Excel spreadsheet with 8 core fields, accessible via SharePoint. The IT lead reviews it quarterly and updates it during the annual HIPAA security review. Staff submit an AI Tool Request Form before using any new AI tool. The documentation identifies a gap: no automated renewal alert capability. The hospital notes this as a future enhancement.	A regional health system documents a SharePoint-based AI registry with automated contract renewal alerts, linked to its vendor management system. Clinical informatics staff audit the inventory semi-annually and cross-reference against IT purchasing records and EHR change logs to identify potential shadow AI. A gap assessment identifies that departmental procurement of AI solutions sometimes bypasses the formal registry process—an escalation procedure is documented.	A large academic medical center documents a purpose-built AI governance platform integrated with its CMDB and procurement systems. Automated scripts query vendor APIs to flag potential new AI system deployments. The inventory is visible to all governance stakeholders via a real-time dashboard. A gap assessment identifies that AI solutions used in research contexts are tracked in a separate research registry with no integration to the clinical registry—a reconciliation process is documented.

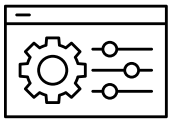

Control 3.7: Identify and document the tools and resources available and needed to carry out validation testing and monitoring for relevant AI models, including responsible AI baseline tests.

A model may perform well on benchmark datasets and poorly on your patient population. Validation testing is the process of independently verifying that an AI system performs as expected—and does not cause harm—in the context of your organization. Similarly, solutions may perform as expected, but degrade in performance or be subject to misuse overtime. Additionally, tracking of relevant key performance indicators related to ROI or outcomes of AI solutions often require longer-term monitoring, and may not be apparent during initial deployment. Monitoring AI solutions can therefore help prevent undesirable shifts in performance and use (which can impact safety or operational integrity), and can help track KPIs and outcomes over time. As always, a risk-based approach to rigor in validation and monitoring is recommended to reduce burden on teams. Additionally, validation and monitoring can be conducted with or without vendor assistance, but the tools and resources must be available.

Testing should go beyond accuracy to include things like usefulness and usability, security/privacy, and clinical workflow integration. At its core, does the solution solve the problem it was intended to solve?

Implementation Guidance

Task 1: Inventory Validation Tools and Expertise

Document what validation and monitoring capabilities your organization has: internal data science or clinical informatics expertise, access to statistical analysis tools, structured testing frameworks (CHAI T&E Frameworks, AHRQ evaluation toolkits), EHR-based performance reporting, and any external validation partners (academic medical center, regional consortium). Note which AI systems have been validated using which tools, and which have not yet been validated. Flag validation gaps—particularly for high-risk clinical AI—as priority remediation items.

Validation and monitoring rigor may vary across different organizations and AI solutions' risk-level. Some organizations consider a desk review of a solution the first step of validation, following it with a small scale pilot or phased deployment where they monitor performance, errors, satisfaction, etc. Others may use retrospective data to test performance before the model ever touches live data streams. Higher risk solutions should be validated and monitored with greater rigor.

Task 2: Define Required Validation Scope and Testing Resources

For each AI system, and depending on risk, document what validation and monitoring should cover. Organizations should consider

(1) Usefulness: Does the tool improve outcomes in your context?

(2) Fairness: Does performance differ across subgroups (race/ethnicity, age, sex, language, insurance status)?

(3) Safety: What are the consequences of incorrect outputs and are fail-safes in place?

(4) Transparency: Can end users appropriately use, trust, or distrust outputs?

(5) Security/Privacy: Does the tool protect PHI?

Document who is responsible for each validation dimension and what data samples or testing environments are needed.

Action: Produce a validation and monitoring requirements document for each AI system specifying: required validation and monitoring dimensions, minimum sample sizes and subgroup requirements, who approves results, and any external resources needed (e.g., third-party validation, vendor data provision).

Task 3: Address the Vendor-Testing Gap and Document Compensating Resources

Many healthcare organizations lack the data science expertise to conduct rigorous validation independently. Document available options to bridge this gap:

(1) Require vendors to provide third-party validation reports and monitoring dashboards

(2) Partner with a nearby academic medical center, health systems, HCCN to share validation and monitoring resources,

(3) Use EHR vendor built-in performance reporting and monitoring tools as a starting point for post-pilot evaluation,

(4) For vendor-managed cloud AI, contractually require performance reporting disaggregated by key subgroups.

(5) Partner with a third-party governance service that provides validation testing for health systems during their RFP process.

6) Partner with vendor to share resources for validation and monitoring

Document which approach is used for each AI system and any gaps that remain.

Real-World Example: Why is Monitoring Important?

One of the issues that can arise when monitoring is not done on an ongoing basis, or when it relies purely on manual checks, is that AI solutions can silently stop working or fail to trigger due to changes in things like data access, hardware updates, etc. One organization shared an example of a "stale application":

"An AI solution for ECG interpretation stopped working after an ECG machine was replaced becaus the model was triggered on hardware metadata. No one noticed for 6 months. The software license continued to be paid, even though the solution was silently inactive and risk predictions stopped reaching the EHR."

Real-World Example: Annual Evaluation Requirements

A large academic medical center shared that their AI governance program requires a mandatory annual evaluation of all deployed AI solutions that includes:

1) Stakeholder check-in to confirm ongoing value

2) Self-reporting of issues or red-flags modeled after an IRB annual review approach

3) Review of usage statistics and outcome variables where applicable.

This type of defined interval review goes beyond ongoing monitoring of performance and bias. Check-ins like this serve as an opportunity to evaluate whether the solution is still in use, having the intended impact, whether a better solution is available, or how the current solution or its use may need to be improved.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small hospital documents that it uses a structured 30-day pilot with two clinical champions before deploying any new clinical AI tool. It documents CHAI's T&E framework checklist as its primary validation resource. A gap assessment for a recently acquired chronic disease risk model identifies that the vendor has not provided disaggregated performance data for its predominantly Latino, uninsured patient population—a request for this data is documented as a procurement condition.	A regional health system documents a formal validation checklist (adapted from CHAI's T&E framework) administered by its clinical informatics team. For a new radiology AI, the team documents: 500-case local validation, three demographic subgroup analyses, radiology chief sign-off required, and AI Governance Committee approval before go-live. A gap assessment notes that validation expertise for NLP-based clinical AI is limited internally—a partnership with a regional academic center is documented as the compensating resource.	A large academic medical center documents a formal AI Validation Protocol with minimum sample sizes, subgroup analyses, and independent reviewer requirements by risk tier. High-risk diagnostic AI undergoes a prospective pilot with formal IRB review. Validation results are published in peer-reviewed literature. A gap assessment identifies that monitoring resources post-deployment are less systematically documented than pre-implementation validation—a monitoring framework is added as a concurrent workstream.

Note: What to Do When Validation Fails

When validation fails, organizations have options:

Vendor/Solution remediation — Document the specific failure and formally request the vendor/developer address it before wider deployment proceeds, with a defined timeline and re-validation criteria.

Conditional deployment with enhanced monitoring — Deploy in a limited, constrained context with explicit safeguards, heightened monitoring, and a defined re-evaluation date. Conditions must be documented and enforced.

Phased pilot with a defined off-ramp — Begin with a structured pilot (single department, limited cohort, fixed time window) with a formal decision gate and pre-defined criteria for halting.

Defer pending additional evidence — Pause deployment while gathering more data, conducting further testing, or waiting for external evidence. Most appropriate when failure reflects insufficient local data rather than a clear performance problem.

Reject the tool — Decline deployment entirely. The right answer when failures involve patient safety, systematic bias, or when a vendor refuses transparency or remediation.

Organizations should define what success and failure thresholds are based on solution risk level and use-case specific considerations.

Computing Resources

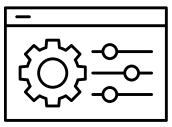

Control 3.8: Identify and document the computing resources available and needed to conduct recurring cybersecurity audits of AI infrastructure; enforce access controls and HIPAA compliance.

AI infrastructure—the hardware, cloud services, software environments, and networks that power AI systems—represents a distinct and critical governance responsibility. Compromising AI infrastructure can impact clinical decision support, patient safety, and operations at critical moments, expose large volumes of training data, or allow adversaries to manipulate model behavior at scale. Before an organization can conduct effective computing infrastructure audits, it must know what infrastructure it has and what resources are needed to audit it systematically.

Implementation Guidance

Task 1: Inventory Your AI Computing Infrastructure

Before you can audit AI computing infrastructure, you must know what it consists of. Maintain a computing resource inventory that documents:

- Physical hardware (servers, GPUs, AI accelerators) used for AI training or inference
- Cloud service dependencies (which AI systems run on AWS, Azure, Google Cloud, or private cloud)
- Software environments and dependencies (libraries, frameworks, containers)
- Network connections between AI systems and clinical data sources, and Third-party API integrations

Many organizations overlook this step, treating AI infrastructure as "just IT infrastructure." AI workloads have distinct characteristics (large data volumes, specialized hardware, API-heavy architectures) that may not be captured by standard IT asset management.

Task 2: Identify Audit Resources and Define Audit Scope

Document what resources are available to audit AI computing infrastructure: internal IT security staff, external security firm, SIEM (Security Information and Event Management) platform, and automated monitoring tools. Define audit scope:

(1) Access control review—who has access to AI models, training data, and inference environments?

(2) Patch and vulnerability management—are AI software dependencies kept up to date?

(3) Backup and recovery—can AI systems be restored after ransomware or system failure? What is the RTO?

(4) Logging and monitoring—are AI system activities logged to enable anomalous behavior detection? Document gaps between required audit scope and available resources.

Task 3: Align Computing Resources with HIPAA and Business Continuity

Document how AI computing infrastructure is incorporated into your broader HIPAA Security Rule compliance program (HIPAA Technical Safeguard requirements apply fully to AI infrastructure). Ensure AI systems are incorporated into Business Continuity and Disaster Recovery (BCDR) plans—the failure of a clinical AI system during an outage should be anticipated, not improvised. Document: which AI systems are covered under vendor Business Associate Agreements, what recovery procedures exist for each AI system, and what computing resources are reserved for AI system restoration.

Real-World Examples

Computing resources, requirements, and related procurement can vary by tool type and organizational context. For example an academic medical center stated that their approach to computing resources varies by tool. For some, the vendor brings a specific server into the HDO's infrastructure. For others, the HDO stands up it's own compute environment. Cybersecurity requirements at this institution require a robust cybersecurity evaluation process.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A critical access hospital documents its AI computing infrastructure as an addendum to its existing annual HIPAA Security Risk Analysis: cloud-hosted AI solutions, their data connections, and access permissions. Access logs are reviewed quarterly by the IT lead. The documentation confirms AI components are covered under vendor HIPAA Business Associate Agreements and that backup procedures include AI-generated data.	A regional health system documents bi-annual cybersecurity audits that include AI infrastructure as a distinct scope area—internal IT security team with external validation every other year. Audit reports go to the CIO and AI Governance Committee with remediation items tracked in a security log. AI system recovery procedures are documented in the BCDR plan. A gap assessment identifies that one cloud AI tool lacks documented RTO/RPO targets—a remediation item is assigned.	A large academic medical center documents continuous SOC-integrated monitoring with AI-specific alerting rules, quarterly tabletop exercises simulating AI infrastructure failure scenarios (including ransomware and vendor outages), and a dedicated AI infrastructure section in its enterprise BCDR plan. A gap assessment identifies that a recently acquired AI tool uses a novel cloud architecture not yet covered by existing monitoring rules—an engineering workstream is created to add coverage.

Human Resources

Control 3.9: Establish an AI workforce competency framework (or related document) that defines: roles and responsibilities of AI governance personnel using AI minimum qualifications and/or certification. Include ongoing training programs to update personnel on governance concepts such as AI safety and bias mitigation.

Governance controls without clear human ownership do not get implemented. Organizations consistently identified roles and responsibilities—both internal and contracted—as the area they struggle with most, including defining competencies for who does what in AI governance. This control asks organizations to build a structured, documented answer to that question: who is responsible for AI governance activities, what qualifications or training do they need, and how does the organization ensure those capabilities are maintained over time?

It is typical that there will be variation in who actually owns activities like vendor AI evaluations or security assessments across different organizations. However, a competency framework and RACI matrix make these accountabilities explicit and sustainable.

Implementation Guidance

Task 1: Define AI Governance Roles and Responsibilities

Document the roles involved in AI governance at your organization and their specific responsibilities. Roles commonly include: AI Governance Committee Chair or lead, Clinical AI Champion(s) (by department or care area), AI Inventory Owner, AI Security Lead (often IT or CISO function), Compliance/Legal representative, and Vendor Management Lead. For each role, document: who currently fills it (by title, not name), their specific governance responsibilities, and whether that role is staffed, shared, or a gap. Consider using a RACI matrix (Responsible, Accountable, Consulted, Informed) to make ownership explicit for each control in this playbook.

Action: Produce a Governance Roles and Responsibilities document (or RACI matrix) mapping each Domain 3 control to a responsible role. This document becomes the accountability anchor for your governance program.

Task 2: Define Minimum Qualifications and Competencies

For each AI governance role, document minimum qualifications or competencies. These need not require formal certification for most roles—but should be explicit. Examples: AI Inventory Owner: familiarity with IT asset management and contract management; AI Security Lead: existing cybersecurity expertise plus AI-specific threat awareness (OWASP AI Security Top 10, MITRE ATLAS); Clinical AI Champion: clinical domain expertise, AI literacy training, CHAI Responsible AI principles; Compliance Representative: HIPAA expertise, working knowledge of FDA AI guidance, awareness of emerging state AI regulations. For organizations subject to state AI regulations (e.g., Texas requiring AI training), document those requirements and how the competency framework addresses them.

Task 3: Establish Ongoing AI Governance Training Programs

AI governance is not a static skill set—the technology, regulatory environment, and risk landscape evolve continuously. Document a plan for ongoing training of AI governance personnel:

(1) Initial onboarding training for new governance personnel

(2) Annual refresher training covering updates to AI safety practices, bias mitigation, and regulatory changes

(3) Role-specific training pathways—clinical champions may need different training than IT security staff

(4) Resources for training: CHAI educational materials, AMA AI guidance, academic CME programs on health AI, and vendor-specific training. Track completion and tie training to governance role renewal.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small FQHC or community hospital documents a simple two-page roles and responsibilities matrix: the IT director owns security controls (Controls 1-3, 8), the CMO or Medical Director owns clinical validation (Control 7) and model card completion (Control 4), and the compliance officer owns documentation and inventory (Controls 5-6). AI governance training consists of annual review of CHAI's Responsible AI Guide and completion of one online AI literacy module, documented in personnel records. The organization notes that all AI governance roles are shared duties, not dedicated positions.	A regional health system develops a formal AI Governance RACI matrix covering all Domain 3 controls and assigning ownership across IT Security, Clinical Informatics, Compliance, Legal, and Operations. Minimum competencies are defined for the AI Inventory Owner and Clinical AI Champion roles. Annual training is built into existing staff development workflows, with CHAI educational materials as the primary resource. The health system tracks training completion in its HR system.	A large academic medical center develops a comprehensive AI Workforce Competency Framework with defined competency levels (awareness, practitioner, expert) for each governance role. It offers formal AI governance training through its CME program and partners with academic programs for advanced training in AI safety and fairness evaluation. A dedicated AI Governance Office leads training coordination. The institution publishes sample job descriptions for AI governance roles as a contribution to the broader CHAI community.

Key Resources & Tools

Sample RACI for Domain 3: Organizational Resources

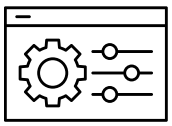

Control 3.10: Identify when clinicians and staff need to be trained on specific AI solutions prior to implementation; also, determine if certain AI solutions require regular training/competency assessments.

AI solutions do not deliver safe, effective care on their own—they depend on clinical users who understand how to interpret their outputs, recognize their limitations, and know when not to rely on them. Workshop participants—including representatives from organizations with active AI governance programs—emphasized that clinician time is constrained, making training requirements that are poorly designed or disproportionate to risk likely to fail. This control asks organizations to develop a structured, risk-tiered approach: identifying which AI solutions require training before use, what that training should cover, and whether ongoing competency assessment is warranted. For further details on this, please refer to the Domain 4.5 Education, Training, and Feedback Playbook.

Implementation Guidance

Task 1: Identify AI Solutions Requiring Pre-Implementation Training

For each AI system, determine whether pre-implementation training is required for clinical users or operational staff. Not all AI solutions require formal training—a scheduling optimization tool may need only basic orientation, while a clinical decision support tool used at the bedside may require formal clinical training and competency demonstration. Document the training requirement for each AI system in the AI inventory (see Control 6) at the time of deployment review. Factors that increase training necessity: direct clinical decision support, outputs that clinicians may over-rely on, novel interfaces or workflows, high-risk patient populations, or evidence of common misinterpretation in literature.

Task 2: Define Training Content and Format by AI Solution

For each AI solution requiring training, document what the training should cover: (1) What the AI does and does not do—intended use and out-of-scope uses, (2) How to interpret the tool's outputs, including confidence scores or probability estimates, (3) Known limitations and failure modes—when to distrust the tool's output, (4) What action clinicians should take if the output appears incorrect or unusual, (5) How to report performance concerns or adverse events related to the AI tool, (6) Applicable regulatory requirements (e.g., documentation requirements for AI-assisted decisions). Where vendor training materials are available, evaluate them against these criteria and supplement where needed.

Task 3: Determine Ongoing Training and Competency Assessment Requirements

Some AI solutions warrant periodic retraining or competency assessment—particularly when the AI model is updated, when clinical evidence about the tool's performance changes, or when the clinical context of use shifts. Document for each AI solution: whether ongoing training is required, at what cadence (annual, following model updates, following significant clinical events), whether formal competency assessment is warranted, and who is responsible for managing training completion. For high-risk clinical AI, consider mirroring the annual check-in model (similar to IRB continuing review) in which ongoing clinical value, user feedback, and usage statistics are reviewed at defined intervals. Additionally training boosters may be required if monitoring reveals misuse or lack of use.

Illustrative Examples by Organization Size

The following scenarios are illustrative only. They represent plausible approaches based on resource level and are not drawn from specific documented organizations or published case studies.

Small Clinic, Critical Access Hospital, or Community Health Center	Mid-Size or Regional HDO	Large System or Academic Medical Center
A small FQHC uses a simple training checklist built into its AI deployment workflow: for each new AI tool, the clinical champion completes a brief structured orientation for all users covering what the tool does, how to read its output, and when to flag concerns. Training completion is noted in personnel records. For its highest-risk tool (a chronic disease risk score), the organization requires annual refresher orientation tied to its quality improvement review cycle.	A regional health system incorporates AI training requirements into its standard clinical education workflow. A clinical informatics analyst documents training requirements for each AI tool in the AI inventory at time of approval. For a new sepsis alert AI, a 30-minute online module is required before access is granted; completion is tracked in the LMS. Ongoing competency is assessed annually via a brief case-based quiz. The system uses vendor training materials supplemented by local clinical context.	A large academic medical center conducts a formal training needs assessment for each AI system as part of the pre-implementation validation process. Training requirements are tiered by risk level and user role. High-risk diagnostic AI (e.g., radiology AI) requires in-person or synchronous training with demonstration of competency before use. The institution runs annual AI literacy workshops for all clinical staff and tracks participation in its CME system. Training data is reviewed by the AI Governance Committee as part of annual AI system review.

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 and Mitigation Approaches	Mitigation
Data Infrastructure Costs	High upfront costs and complexity of upgrading legacy systems to support AI; particularly acute for smaller organizations.	Start with EHR-native or vendor-managed cloud AI to leverage existing infrastructure before investing in custom infrastructure.
Data Security Complexity	Maintaining security standards across AI systems is complex, particularly for generative AI that may output unexpected content containing PHI.	Require vendors to document their security architecture and provide SOC 2 reports; apply output filtering for generative AI.
Security Assessment Restrictions	Vendors frequently restrict or decline permission for penetration testing and vulnerability scans.	Include testing rights in contract language at procurement stage—retrofitting contracts is significantly harder.
Shadow AI	Staff use of unauthorized AI solutions (public generative AI, consumer tools) is widespread and difficult to detect.	Focus AI policy dissemination on making the governance process easy and visible; create a simple, low-friction intake pathway so staff want to use the official process.
Staffing and Expertise Gaps	Shortage of staff with both technical and governance expertise; particularly acute in smaller organizations and for specialized skills like AI security and bias evaluation.	Leverage shared services models, regional health system consortia, and primary care/health center controlled networks. Working together with regional organizations who are interested in similar solutions can help reduce or share burden. Also sharing learnings across trade and other membership organizations can help.
Immature AI Cybersecurity Tools	Commercial tools for AI-specific security testing (adversarial robustness, model theft detection) are less mature than traditional IT security tools.	Use existing frameworks to structure evaluations even where dedicated tools don't yet exist.
Computing Resource Documentation	AI computing infrastructure is frequently underdocumented or not tracked separately from general IT assets.	Start with a simple AI infrastructure addendum to existing IT asset management processes.
Documentation Burden	Data scientists and clinical staff resist detailed documentation requirements	Integrate documentation into existing workflows (e.g., auto-populate model

	as burdensome.	card fields from MLOps platforms); start with minimum viable fields and expand over time
Outdated Documentation	Vendors and developers may not always communicate updates to their AI systems, rendering model cards and registry information as outdated.	Make notification of material changes a contractual obligation at procurement, rather than relying on vendor goodwill. Specifically, AI vendor contracts and BAAs should include a defined notification SLA, requiring the vendor to notify the organization within a specified timeframe whenever a material change occurs to the AI system or whenever a new version is deployed. Material changes should be explicitly defined in the contract to include changes, for example, to intended use, scope of use, or out-of-scope uses, added features, known performance changes or degradations, and changes to underlying infrastructure or subprocessors. Because many AI solutions are updated frequently, it is often not feasible to update organizations anytime training data is updated, or model changes are made.

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

This document is licensed under a Creative Commons Attribution-Non-Commercial-No Derivatives 4.0 International License (CC BY-NC-ND 4.0).

You are free to share this material (copy and redistribute it in any medium or format) under the following terms:
Attribution: You must give appropriate credit, provide a link to the license, and indicate if changes were made.
Noncommercial: You may not use the material for commercial purposes.
No Derivatives: If you remix, transform, or build upon the material, you may not distribute the modified material.

For more information about this license, visit creativecommons.org/licenses/by-nc-nd/4.0/ .