AI Governance Playbooks: Introduction & Background

Guiding responsible AI adoption

Why AI Governance?

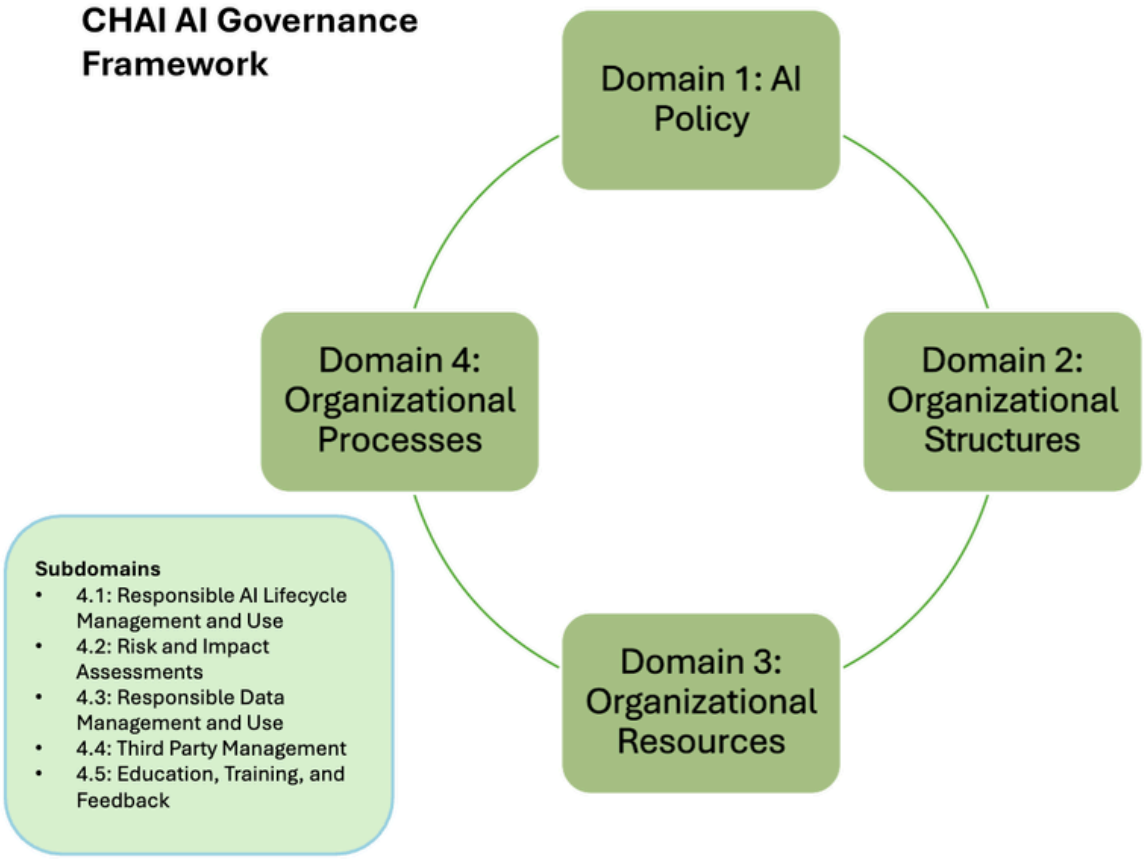
Effective AI governance enables healthcare organizations to make consistent, accountable decisions about AI use while protecting patient safety, workforce trust, and regulatory compliance.

CHAI AI Governance Framework

CHAI's initial Governance Framework was derived from the principles and structure of *ISO/IEC 42001: Artificial Intelligence Management Systems*. This foundation provided a systematic approach to AI governance, emphasizing key governance domains. The framework was then adapted to healthcare delivery organizations to account for sector-specific considerations.

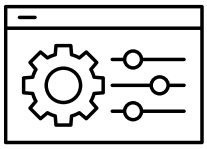
CHAI's Governance Framework consists of four core domains (AI Policy, Organizational Structures, Organizational Resources, and Organizational Processes), with five subdomains under Organizational Processes. All playbooks in this series are designed for healthcare delivery leaders from organizations of all sizes, who want to build or expand on capabilities and governance focused on the implementation and use of AI within their organizations.

Note: AI governance controls mentioned here may also apply to other non-provider organizations, given that the development, evaluation, and use of AI in these organizations may also impact clinical decision-making, treatment access, care quality, or patient outcomes. Implementation guidance described here focuses on healthcare delivery organizations, but can be adapted accordingly for different settings.



How to Use These Playbooks

Each Playbook includes **baseline operational controls**, which were consensus-defined (see later sections for full methodology) by healthcare delivery organizations as "must-haves" for health AI governance. These controls will be accompanied by the following symbol:



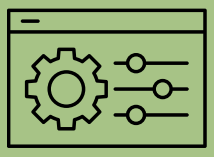
The **implementation guidance** that accompanies the governance controls are **recommendations** based on known practices. It is assumed and appropriate that implementation will differ based on organizational goals, resources, and maturity. To demonstrate this flexibility, we have provided examples from different organizational contexts and tools to help support organizations through the design and implementation of their unique governance journeys. Implementation guidance will be preceded by the following symbol:



Guiding Note

It is important to note that some of the recommended governance controls in the CHAI AI Governance Framework may not require entirely new processes, but rather the adaptation or expansion of existing structures already in place within healthcare delivery organizations. While the CHAI AI Governance Framework presents controls in isolation for clarity and consistency, its intent is not to imply that each control necessitates a net-new implementation. For example, if a health system already maintains a process for reporting data breaches or privacy violations, that same reporting infrastructure could be leveraged to address issues arising from AI implementations—simply by incorporating a new category or tagging mechanism for AI-related concerns.

Organizations should adopt a risk-based framework across all of the controls, with higher-risk AI solutions requiring more in depth processes (e.g. intake, evaluation, validation, data management, life cycle management, monitoring, etc.) and documentation, and lower-risk AI solutions following more simplified processes and documentation.



Baseline Operational Controls for All Domains

This is an overview of all baseline operational controls across the AI governance domains and subdomains.

Control ID	Control Description
AI Policy 1.1	Develop a formal AI Policy document, approved by organizational leadership, and make accessible to all relevant staff.
AI Policy 1.2	Define key concepts in the AI Policy (e.g., definition of AI, scope of AI solution use, AI approval process).
AI Policy 1.3	Establish an appropriate individual or individuals with designated oversight for policy changes.
Organizational Structures 2.1	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.
Organizational Structures 2.2	Establish an AI Governance Committee - a formal committee structure bringing together all governance stakeholders and responsibilities as defined by the organization.
Organizational Structures 2.3	Establish an intake and monitoring process for thorough investigation of prospective solutions and mechanisms for incident detection, feedback, and follow-up.
Organizational Structures 2.4	Establish protocols for escalating and communicating concerns/incidents for resolution to the governance structure.
Organizational Resources 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.
Organizational Resources 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.
Organizational Resources 3.3	Maintain comprehensive records of AI-related assets including datasets, models, and tools to ensure traceability.
Organizational Resources 3.4	Identify and document tools, resources, and methods available and needed to implement AI cybersecurity evaluation (penetration testing, adversarial robustness, etc.).
Organizational Resources 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.

Organizational Resources 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.
Organizational Resources 3.7	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.
Organizational Resources 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.
Organizational Resources 3.9	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.
Organizational Resources 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.
Responsible AI Lifecycle Management and Use 4.1.1	Develop and maintain a process that appropriately assesses and documents intended use of each AI solution.
Responsible AI Lifecycle Management and Use 4.1.2	Define responsible AI principles, as aligned with the organization's mission and goal (e.g., usefulness, fairness, safety, transparency, security, privacy); document the objectives and processes that the organization will use to align with these principles across relevant lifecycle steps.
Responsible AI Lifecycle Management and Use 4.1.3	Operationalize a process to evaluate defined responsible AI principles (e.g., usefulness, fairness, safety, transparency, security, privacy) and reference industry recommended evaluation frameworks where available.
Responsible AI Lifecycle Management and Use 4.1.4	Define a process to monitor the safety, performance, and responsible use of AI solutions.
Responsible AI Lifecycle Management and Use 4.1.5	Establish a monitoring mechanism or process for ongoing performance reporting; monitoring mechanisms or processes should flag unexpected patterns for review by responsible parties.

Responsible AI Lifecycle Management and Use 4.1.6	Establish a process to inform relevant stakeholders about responsible AI principles and processes (e.g., definitions, evaluation, monitoring).
Risk and Impact Assessments 4.2.1	Conduct a Risk Categorization assessment (pre-deployment) to categorize risk in the context of the organization's AI use case scope.
Risk and Impact Assessments 4.2.2	Document how Risk Categorization is evaluated, recorded, retained, and reviewed.
Risk and Impact Assessments 4.2.3	Leverage a rigorous Risk Assessment to evaluate AI solutions that are categorized as higher risk (as defined by the organization).
Risk and Impact Assessments 4.2.4	Document how Risk Assessment is evaluated, recorded, retained, and reviewed.
Risk and Impact Assessments 4.2.5	Conduct and document an AI System Impact Assessment that includes evaluation of risks and benefits for AI solutions.
Responsible Data Management and Use 4.3.1	Establish and maintain a structured data acquisition register for data used in AI systems.
Responsible Data Management and Use 4.3.2	Establish comprehensive agreements defining terms for permissible/non-permissible uses, re-identification and redistribution preferences, requirements for vendor security compliance, audit rights, and data return/destruction terms.
Responsible Data Management and Use 4.3.3	Establish written HIPAA-compliant de-identification policies; evaluate PHI datasets to select the appropriate method; validate de-identification before AI training use; maintain oversight of linked datasets.
Responsible Data Management and Use 4.3.4	Establish a process and owner for recurring AI data quality and bias evaluation to help identify quality gaps and known sources of bias.
Responsible Data Management and Use 4.3.5	Establish a standard operating procedure (SOP) for AI data preparation including rules for data cleaning, encoding practices, normalization/standardization methods, and partitioning procedures.
Third Party Management 4.4.1	Outline requirements for disclosure of known model limitations and risks; define roles and responsibilities pertinent to the AI solution and known risks; include contractual and operational risk mitigation provisions to de-risk the use of the solution.

Education, Training, and Feedback 4.5.1	Develop a process to ensure all relevant staff receive training on AI solutions, including their purpose, limitations, and related information. Reinforce training via recurring briefings and optional seminars for deeper engagement.
Education, Training, and Feedback 4.5.2	Provide all patients, no later than the date of first service delivery, with a comprehensive Notice of Privacy Practices (NPP), including how protected health information (PHI) may be used by AI solutions, in compliance with HIPAA.
Education, Training, and Feedback 4.5.3	Develop a process for patient transparency and/or consent where appropriate for relevant AI solutions.
Education, Training, and Feedback 4.5.4	Develop a process to document, escalate, and notify relevant parties about AI incidents. Ensure it is clearly communicated that all parties involved in an AI incident share responsibility for reporting any voluntary or mandatory incidents.
Education, Training, and Feedback 4.5.5	Ensure information included in AI incident communication falls under relevant federal and state laws and protections.
Education, Training, and Feedback 4.5.6	Provide whistleblower protections for both internal and external reporting, including anonymity, confidentiality, non-retaliation assurances, and clear disclosure of data retention policies.

Playbooks in This Series

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

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 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.

Phase 1 Methodology

The initial Phase 1 development of CHAI's Governance Framework followed a structured, multi-phase process designed to ensure both rigor and relevance to the healthcare delivery context. The process combined international standards alignment, expert validation, and broad stakeholder engagement to define a set of baseline operational controls for healthcare delivery organizations.

1. Foundation and Adaptation

CHAI's initial Governance Framework was derived from the principles and structure of ISO/IEC 42001: Artificial Intelligence Management Systems. This foundation provided a systematic approach to AI governance, emphasizing key governance domains. The framework was then adapted to healthcare delivery organizations to account for sector-specific considerations.

2. Evidence Review and Expert Input

To inform the initial draft, CHAI conducted a targeted review of health AI governance frameworks, standards, and implementation practices across health systems and technology vendors. This was supplemented by a series of one-on-one interviews with subject matter experts, including clinicians, data scientists, compliance officers, and legal professionals, to contextualize governance principles within real-world healthcare workflows.

3. Iterative Stakeholder Engagement

Draft versions of the Governance Framework were shared with CHAI's Health System Advisory Board (HSAB) for iterative review and refinement. The resulting content formed the basis of a structured series of governance workshops, designed to solicit feedback across diverse perspectives. We prioritized feedback from healthcare delivery organizations of various sizes, given that they will be the primary user and audience of these playbooks. We additionally engaged governance platform developers, health technology vendors, and legal representatives from both provider and vendor organizations.

4. Synthesis of Workshop and Survey Feedback

Feedback collected during workshops was distilled and analyzed thematically to identify areas of convergence and divergence across stakeholder groups. To include organizations unable to participate in live sessions, CHAI distributed an asynchronous survey mirroring the workshop content. This ensured comprehensive representation and input from the broader CHAI community.

5. Consensus-Based Baseline Operational Control Selection Criteria

Controls were evaluated and prioritized using predefined inclusion thresholds based on aggregated feedback:

- Primary - Baseline Operational Control: ⩾67% of both Governance Platform and Healthcare Delivery respondents identified the control as a "Must Have".
- Secondary - Baseline Operational Control: ⩾50% and <67% of either Governance Platform or Healthcare Delivery respondents identified the control as a "Must Have".

- Tertiary - Baseline Operational Control: <50% of both Governance Platform and Health System respondents identified the control as a “Must Have”.

Overall, ~200 individual organizations were invited to and participated across workshops and surveys, representing primarily a mix of healthcare delivery organizations and governance platforms, with a few perspectives from health technology vendors, and legal stakeholders.

Workshop Statistics & Attendance

CHAI invited over 270 members and subject matter experts across 215 unique organizations.

The below statistics reflect the total number of unique participants across the workshops.

Workshop 1

Registered: 121; Attended: 89

- ⩽10 AI Solutions Deployed: 28%
- 11+ AI Solutions Deployed: 46%
- Governance Platform: 26%

Workshop 2

Registered: 125; Attended: 80

- ⩽10 AI Solutions Deployed: 12%
- 11+ AI Solutions Deployed: 58%
- Governance Platform: 30%

Workshop 3

Registered: 121 ; Attended: 45

- ⩽10 AI Solutions Deployed: 15%
- 11+ AI Solutions Deployed: 40%
- Governance Platform: 45%

Survey Participants

CHAI also included responses from the identical, asynchronous survey.

Responses: 15

Unique Organizations: 10

- ⩽10 AI Solutions Deployed: 20%
- 11+ AI Solutions Deployed: 40%
- Governance Platform: 40%

Phase 2 Methodology

Phase 2: Engaged subject matter experts and members to build out governance controls, implementation guides, including practical steps grounded in context, tools and resources, and examples of how differently resourced and sized organizations may align.

- Gathered 1:1 insights through interviews with leaders at healthcare delivery organizations
- Member & Advisory Group Feedback Period on output of Governance Workshops (Controls)
- Solicit Interest in Governance Playbook Workgroups (November 2025-Jan 2026)
- Convened Governance Workgroup & Draft v0.1 & v0.2 of Governance Playbooks (Jan 2025-April 2026)
  - Workgroup members conducted two full draft reviews with comments and discussions for each playbook
- 3-Week Member, Board, Workgroup, and Advisory Group Feedback Period
- Feedback Prioritization & Resolution
- Output: Version 1.0 of Governance Playbooks

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 explicitly opted-in to being named 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

Glossary

Term	Definition
AI inventory	The organization-wide registry of all deployed AI systems, including operational, clinical, administrative, and research AI. The AI inventory is the foundational operational tool for AI governance; it captures intended use, risk tier, monitoring owner, contract status, and links to model cards (see Domain 3, Control 6).
AI model	A trained computational artifact — comprising parameters, weights, and architecture — that produces outputs from inputs. The model is the core technical component of an AI system but is not itself a complete AI system.
AI solution	An AI system deployed to address a specific organizational use case, including the model, the workflow integration, user interface, and any human-in-the-loop processes. Used in CHAI's Applied Model Card context.
AI system	The complete deployed assembly of one or more AI models with associated technical infrastructure, data pipelines, integration logic, and human workflows. The primary unit of analysis for governance under NIST AI RMF, ISO/IEC 42001, and the CHAI framework.
AI Governance Committee (AIGC)	The formal organizational body responsible for AI intake review, risk categorization, lifecycle oversight, and incident escalation. May be a dedicated committee, a subgroup of an existing governance body (e.g., Data Governance Council), or a distributed structure of named individuals (see Domain 2).
Agentic AI	AI systems that take autonomous actions across multiple tools, systems, or environments to accomplish a task, rather than producing recommendations for a human to act on. Agentic AI introduces distinct governance considerations relating to scope of autonomy, action reversibility, tool access, persistent memory, and audit-trail granularity.
Applied Model Card (AMC)	CHAI's standardized model card template, designed for healthcare delivery contexts. Documents intended use, performance, training data characteristics, known limitations, and Responsible AI principles for a specific AI solution in context.
Bias	Systematic difference in the performance, training data composition, or outputs of an AI system across groups or contexts. Bias may originate in data (representation, label, temporal, survivorship, selection, measurement bias) or in algorithmic choices. Bias is a property of the AI system; it is distinct from disparity (a property of population outcomes).
Business Associate Agreement (BAA)	A contract required under HIPAA between a covered entity and a business associate that creates, receives, maintains, or transmits Protected Health Information on the covered entity's behalf. For AI vendors, BAAs require AI-specific

	provisions covering model training restrictions, derivative-data rights, and return/destruction of model artifacts.
Centralized AI governance structure	An AI governance model in which decision-making authority, standards, and oversight are concentrated within a single enterprise-level body.
Clinician	Any licensed clinical professional, including physicians, advanced practice providers, registered nurses, pharmacists, therapists, and other licensed care providers. Use “physician” only when the guidance is specific to medical doctors.
Community Health Center (CHC)	A community-based, patient-directed organization that delivers comprehensive, culturally competent, high-quality primary health care services to medically underserved populations. Includes Federally Qualified Health Centers (FQHCs).
Data Use Agreement (DUA)	A HIPAA-required contract between a covered entity and a recipient of a Limited Data Set, specifying permitted uses, recipient obligations, and prohibitions on re-identification.
De-identified data	Data with all 18 HIPAA Safe Harbor identifiers removed (45 C.F.R. § 164.514(b)) or statistically de-identified via Expert Determination (45 C.F.R. § 164.514(c)).
Decentralized AI governance structure	An AI governance model in which primary responsibility for AI oversight is held within individual departments, service lines, or business units with minimal central coordination.
Decision Gate	A documented governance checkpoint within an AI lifecycle stage at which approval status, evidence reviewed, and next review trigger dates are formally recorded. Decision gates mitigate informal progression through lifecycle stages.
Disparity	A documented difference in health outcomes, access, or experience among populations. Distinct from bias: disparity describes outcomes; bias describes a property of the AI system or data.
Distributed AI governance structure	An AI governance model that combines central standards and guardrails with local execution and accountability. Distributed governance often relies on shared templates, risk taxonomies, and escalation pathways.
Evaluation	The measurement of an AI system’s performance against defined criteria. Evaluation produces the evidence used for verification, validation, and ongoing monitoring decisions.
Fairness	A formal property of an AI system relating to the consistency of performance across groups or contexts. One of CHAI’s named Responsible AI principles. Operationalized via metrics applied during testing and monitoring.

Federally Qualified Health Center (FQHC)	A federally designated organization providing primary care services in underserved areas, receiving funding under Section 330 of the Public Health Service Act.
Healthcare Delivery Organization (HDO)	Any organization providing direct patient care, including hospitals, integrated health systems, ambulatory care providers, community health centers, federally qualified health centers, and critical access hospitals. CHAI's primary audience for the AI Governance Playbook series.
Health Center Controlled Network (HCCN)	A federally funded organization that supports member health centers (typically FQHCs) by providing shared technology, data, and quality improvement services.
Impact Assessment (AI System)	A structured evaluation of the broader, system-level consequences of deploying an AI solution, including effects on clinical workflows, patient experience, equity, financial incentives, and organizational culture (see Subdomain 4.2, Control 5).
Ingestion Modality	The pathway through which an AI solution enters a healthcare delivery organization. CHAI defines four primary modalities: (1) Research Models/Solutions, (2) In-house Developed Models/Solutions, (3) Vendor or Co-developed Models/Solutions, and (4) AI-Enhanced Features in Existing Health IT Systems.
Intended use	The defined purpose, user population, clinical or operational context, and scope of an AI solution. Distinct from "intended users" (who uses it) and "permissible uses" (what uses are explicitly allowed).
Limited Data Set (LDS)	Partially de-identified data under HIPAA: direct identifiers are removed, but geographic data, dates, and ages may remain. Requires a DUA before disclosure.
Model card	A structured documentation artifact for an AI system describing intended use, performance, training data characteristics, known limitations, and monitoring information. Used generically; CHAI's specific template is the Applied Model Card (AMC).
Monitoring	Ongoing lifecycle activity focused on evaluating safety, performance, and responsible use of an AI system over time. Includes tracking model drift, override rates, subgroup performance, and adherence to intended use.
Non-medical Drivers of Health (NMDoH)	Non-clinical factors affecting health outcomes, including socioeconomic status, housing, food security, transportation, education, and social support. Also referenced in industry as Social Determinants of Health (SDoH); CHAI's preferred term is NMDoH.
Notice of Privacy Practices (NPP)	The HIPAA-mandated notice provided to patients describing how their Protected Health Information may be used or disclosed.

Patient	The default term for an individual receiving clinical care. Use “patient and family,” “patient and caregiver,” or “patient and community” when the broader unit is intended.
Pilot	A controlled, limited-scale deployment of an AI solution designed to validate real-world performance, usability, workflow integration, and risk assumptions prior to full rollout.
Pre-deployment	The lifecycle stage between design/development and pilot, during which readiness for real-world use is assessed and risks proactively addressed.
Primary Care Association (PCA)	A state or regional organization supporting community health centers, often providing shared infrastructure, training, and contract negotiation support.
Protected Health Information (PHI)	Individually identifiable health information protected under HIPAA, including patient names, dates, contact information, diagnoses, treatments, and any data that could identify an individual.
RACI matrix	A responsibility-assignment table identifying who is Responsible, Accountable, Consulted, and Informed for each governance activity.
Responsible AI (RAI)	CHAI’s set of named principles for the development, deployment, and monitoring of AI systems: usefulness, usability, efficacy, fairness, bias management, safety, reliability, transparency, privacy, and security. When used as an adjective modifier of a CHAI principle or program, capitalize (Responsible AI Guide); when generic, lowercase (responsible AI development practices).
Risk Assessment	A detailed evaluation of the probability and magnitude of harm associated with a higher-risk AI solution, including risk modifiers, mitigation controls, and residual risk acceptance.
Risk Categorization	A pre-deployment triage exercise that classifies an AI solution as Low, Medium, or High risk based on risk domains relevant to the intended use case.
Risk-based approach	An overarching principle in the CHAI framework requiring that the rigor of governance activities (validation, monitoring, documentation, training) scale to the risk of the AI solution.
Shadow AI	AI systems, features, or capabilities developed, procured, embedded, or used by staff without going through the organization’s established AI intake, risk assessment, and approval processes.
Software as a Medical Device	Software intended to be used for one or more medical purposes that performs these purposes without being part of a hardware medical device. SaMD is subject

(SaMD)	to FDA regulation under 21 CFR.
Subdomain	The CHAI framework's second-level structural unit. Domain 4 (Organizational Processes) contains five subdomains: 4.1 Responsible AI Lifecycle Management; 4.2 Risk and Impact Assessments; 4.3 Responsible Data Management and Use; 4.4 Third Party Management; 4.5 Education, Training, and Feedback.
Termination (AI lifecycle)	Planned or unplanned retirement, removal, or disabling of an AI solution, including data disposition, stakeholder notification, and lessons-learned review.
Testing	The discrete procedural activity used to gather evaluation evidence. Includes penetration testing, adversarial robustness testing, A/B testing, silent testing, and prospective deployment testing.
Testing & Evaluation (T&E) Framework	CHAI's consensus-defined set of methods, metrics, measures, and benchmarks for evaluating the responsible use of health AI solutions for a given use case.
Third party (noun) / third-party (adjective)	An organization separate from the healthcare delivery organization that develops, supplies, hosts, or maintains an AI system. Hyphenate as a compound adjective.
Validation	Confirmation that an AI system meets the needs of its intended use in the deploying organization's population and workflow. Distinct from verification (confirmation that the system has been built to specification).
Verification	Confirmation that an AI system has been built to specification — that it does what its documentation states. Distinct from validation.

Acronym & Abbreviations List

Acronym	Full term
AAA	AI & Automation Advisory (organization-specific example)
AAMC	Association of American Medical Colleges
AHA	American Hospital Association
AI*	Artificial Intelligence
aiAE	AI-induced Adverse Event
AIGC	AI Governance Committee
AIGP	Artificial Intelligence Governance Professional (IAPP credential)
AMA	American Medical Association
AMC	CHAI Applied Model Card
API	Application Programming Interface
BAA	Business Associate Agreement
BCDR	Business Continuity and Disaster Recovery
CAH	Critical Access Hospital
CCPA / CPRA	California Consumer Privacy Act / California Privacy Rights Act
CDS	Clinical Decision Support
CfC / CoP	Conditions for Coverage / Conditions of Participation (CMS)
CHAI	Coalition for Health AI
CHC	Community Health Center
CIO	Chief Information Officer
CISA	Cybersecurity and Infrastructure Security Agency

CISO	Chief Information Security Officer
CME	Continuing Medical Education
CMIA	California Confidentiality of Medical Information Act
CMDB	Configuration Management Database
CMO	Chief Medical Officer
CMS	Centers for Medicare & Medicaid Services
CNE	Chief Nursing Executive
CNO	Chief Nursing Officer
CONOPS	Concept of Operations
CPA	Colorado Privacy Act
CSF	Cybersecurity Framework (NIST)
DCAM	Data Management Capability Assessment Model (EDM Council)
DHT	Digital Health Technology
DICOM	Digital Imaging and Communications in Medicine
DID	De-identified Data
DMBOK	Data Management Body of Knowledge (DAMA International)
DMM	Data Management Maturity (CMMI Institute)
DPIA	Data Protection Impact Assessment
DUA	Data Use Agreement
EHR*	Electronic Health Record
ERM	Enterprise Risk Management
EU AI Act	European Union Artificial Intelligence Act

FAIR	Findable, Accessible, Interoperable, Reusable (data principles)
FDA*	U.S. Food and Drug Administration
FMEA	Failure Modes and Effects Analysis
FQHC	Federally Qualified Health Center
FURM	Fair, Useful, Reliable Models (Stanford framework)
GRC	Governance, Risk, and Compliance
GenAI	Generative AI
GPO	Group Purchasing Organization
HCA	HCA Healthcare (organization-specific reference)
HCCN	Health Center Controlled Network
HDO	Healthcare Delivery Organization
HHS	U.S. Department of Health and Human Services
HICP	Health Industry Cybersecurity Practices (HHS 405(d))
HIE	Health Information Exchange
HIPAA*	Health Insurance Portability and Accountability Act
HITRUST	HITRUST Common Security Framework
HSAB	Health System Advisory Board (CHAI)
HSCC	Health Sector Coordinating Council
IAPP	International Association of Privacy Professionals
IMDRF	International Medical Device Regulators Forum
IRB	Institutional Review Board
ISO	International Organization for Standardization

ISO/IEC 42001	ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system
ISO/IEC 5338	ISO/IEC 5338:2023 — AI system life cycle processes
IT*	Information Technology
KPI	Key Performance Indicator
LDS	Limited Data Set
LLM	Large Language Model
LMS	Learning Management System
MCARE	Medical Care Availability and Reduction of Error Act (Pennsylvania)
MDR	FDA Mandatory Device Reporting
MHMDA	My Health My Data Act (Washington)
ML	Machine Learning
MLOps	Machine Learning Operations
MSSP	Managed Security Service Provider
NACHC	National Association of Community Health Centers
NMDoH	Non-medical Drivers of Health
NIST	National Institute of Standards and Technology
NIST AI RMF	NIST AI Risk Management Framework (NIST AI 100-1)
NLP	Natural Language Processing
NPP	Notice of Privacy Practices
NVD	National Vulnerability Database (NIST)
NYPORTS	New York Patient Occurrence Reporting and Tracking System

OCR	Office for Civil Rights (HHS)
ONC	Office of the National Coordinator for Health Information Technology
ONC HTI-1	ONC Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing Final Rule
OWASP	Open Worldwide Application Security Project
PCA	Primary Care Association
PFAC	Patient and Family Advisory Council
PHI	Protected Health Information
PIA	Privacy Impact Assessment
PII	Personally Identifiable Information
PSO	Patient Safety Organization
RACI	Responsible, Accountable, Consulted, Informed
RAG	Retrieval-Augmented Generation
RAI	Responsible AI
RAIG	CHAI Responsible AI Guide
RAIL	Responsible AI Lifecycle (Stanford example)
RCM	Revenue Cycle Management
RFP	Request for Proposal
RN	Registered Nurse
ROI	Return on Investment
RPO	Recovery Point Objective
RRT	Rapid Response Team

RTO	Recovery Time Objective
SaMD	Software as a Medical Device
SLA	Service Level Agreement
SDoH	Social Determinants of Health
SIEM	Security Information and Event Management
SOC	Service Organization Control (SOC 2 Type II)
SOP	Standard Operating Procedure
SOW	Statement of Work
SRA	Security Risk Assessment
SSN	Social Security Number
T&E	Testing and Evaluation (CHAI Framework)
TEFCA	Trusted Exchange Framework and Common Agreement
TGH	Tampa General Hospital (illustrative reference)
TLS	Transport Layer Security
U.S.*	United States
UDS+	Uniform Data System Plus (HRSA)
VOI	Value on Investment
VQT	Virtual Quality Team (illustrative example)
WHO	World Health Organization

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/ .