CHAI Governance Framework for Healthcare Delivery Organizations – Baseline Operational Controls

March 2026 – DRAFT

Background: CHAI is developing a governance framework and set of baseline operational controls that align closely with key domains within CHAI's governance framework. The primary audience for this governance framework is healthcare delivery organizations of all types and sizes. CHAI's governance framework is organized into four overarching domains:

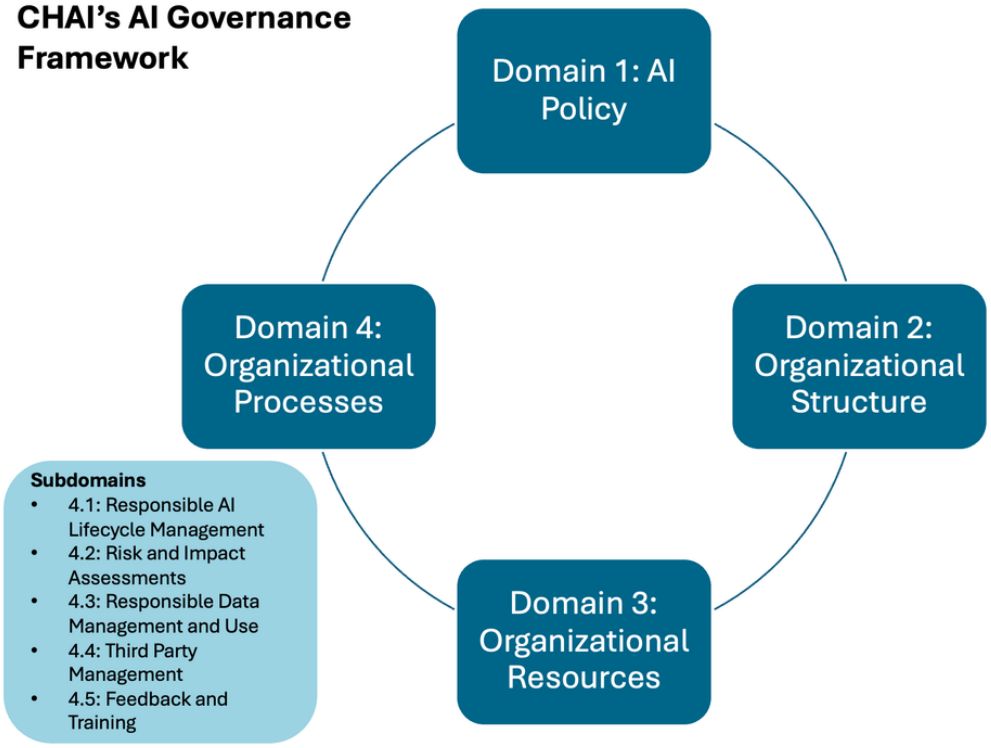

Included below is additional context for the governance framework's domains and subdomains.

Domain 1: AI Policy

Healthcare delivery organizations should establish a formal, organization-wide policy that governs the safe and ethical development, purchasing, and/or use of AI solutions, as appropriate to their organizational context.

Baseline Operational Controls

(1) Develop a formal AI Policy document, approved by organizational leadership, and make accessible to all relevant staff (e.g., clinical staff, IT, compliance officers).

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

(3) Establish an appropriate individual or individuals with designated oversight for policy changes.

Domain 2: Organizational Structure

Based on the context of your organization, establish a cross-functional governance structure with clear roles and responsibilities for overseeing the development, procurement, implementation, use, and monitoring of health AI solutions.

Baseline Operational Controls

(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.

(2) Establish an AI Governance Committee - a formal committee structure bringing together all governance stakeholders and responsibilities as defined by the organization.

(3) Establish an intake process for thorough investigation of prospective solutions and mechanisms for feedback and follow-up.

(4) Establish protocols for communicating concern/incident resolution to stakeholders.

Domain 3: Organizational Resources

Healthcare delivery organizations should account for the availability of resources (e.g., data, tooling, computing, human) required by AI systems and associated assets.

Baseline Operational Controls

(1) Data Resources: Maintain security infrastructure (aligning with HITRUST, SOC 2, NIST CSF, CISA best practices) to protect against unauthorized access, disclosure, or re-identification of data; encrypt data at rest and in transit.

(2) Data Resources: Conduct regular security assessments and vulnerability scans to identify and address potential security risks and mitigate known vulnerabilities and risks.

(3) Tooling Resources: Adopt a Model Card or other standard documentation framework for all AI systems, mandating completion of fields such as model ID, intended use, primary end user, and performance metrics – integrated into MLOps or AI inventory registry.

(4) Tooling Resources: Implement and maintain an AI inventory registry, auditing all internal and third-party models for gaps, capturing model details, vendor contact info, and usage restrictions, with (semi-)automation to monitor renewal dates and policy compliance.

(5) Tooling Resources: Carry out pre-implementation validation testing for every AI model, including baseline validation tests for responsible AI principles such as Usefulness, Fairness, and Safety.

(6) Tooling Resources: Implement mechanisms to conduct AI cybersecurity evaluation for penetration testing, adversarial robustness assessments, API security, data encryption, and model theft protections (as examples).

(7) Computing Resources: Conduct recurring cybersecurity audits of AI infrastructure and operations to identify vulnerabilities that may be exploited by adversaries; enforce access controls and compliance with HIPAA; ensure systems can recover from disruptions.

(8) Resource Documentation: Maintain comprehensive records of AI-related assets including datasets, models, and tools to ensure traceability.

Domain 4: Organizational Processes

Operationalize the processes that aid in AI solution and system management across the lifecycle, including how to align with responsible AI principles. Included below is additional context for each subdomain within domain 4.

Subdomain 4.1: Responsible AI Lifecycle Management

Healthcare delivery organizations should implement processes that support responsible lifecycle management and use of AI systems while identifying and documenting broader objectives.

Baseline Operational Controls

(1) Develop and maintain a process that appropriately assesses and documents intended use of each AI system.

(2) Operationalize responsible AI principles, as defined by the organization; document the objectives and processes that the organization will use to align with these principles across relevant lifecycle stages.

(3) Objectively evaluate responsible AI principles (such as usefulness, fairness, safety, transparency, security, privacy) and reference industry recommended evaluation frameworks.

(4) Implement a process to monitor and regularly evaluate the safety performance and responsible use of AI-enabled solutions.

(5) Establish a monitoring mechanism for frequent reporting on performance via a dashboard or similar interface. Monitoring mechanism should flag unexpected usage patterns for manual review by IT and compliance staff.

(6) Integrate Responsible AI Use policy training into new hire orientation and annual compliance training for clinicians, administrators, IT staff, and all relevant stakeholders.

Subdomain 4.2: Risk and Impact Assessments

Healthcare delivery organizations should document processes used to conduct AI risk and impact assessments.

Baseline Operational Controls

(1) Conduct a Risk Categorization assessment (pre-deployment) to categorize risk in the context of the organization's AI use case scope. This is a high-level categorization covering different dimensions of risk.

(2) Document how Risk Categorization is evaluated, recorded, retained, and reviewed.

(3) Leverage a rigorous Risk Assessment to evaluate AI solutions that are categorized as higher risk (as defined by the organization).

(4) Document how Risk Assessment is evaluated, recorded, retained, and reviewed.

(5) Conduct and document an AI System Impact Assessment that includes evaluation of risks and benefits for all AI solutions.

Subdomain 4.3: Responsible Data Management and Use

Healthcare delivery organizations should identify and document processes associated with data management and use as it relates to AI systems. Many healthcare delivery organizations have preexisting documented processes and procedures related to data quality, use, privacy, security, and management. Organizations should update existing processes to account for AI system-specific requirements.

Baseline Operational Controls

(1) Data Use & Enhancement: Contractually prohibit unauthorized re-identification, linking, or redistribution of de-identified data, especially when used to train or fine-tune AI models, via robust Data Use Agreements; Ensure agreements specify return or destruction of data after use.

(2) Data Use & Enhancement: Clearly define permissible AI use cases in Data Use Agreements and prohibit all other use.

(3) Data Use & Enhancement: Ensure all third-party vendors comply with the healthcare organization's data security and privacy requirements, as outlined in the Data Use Agreement. These requirements should include, at a minimum, encryption, access controls, and regular security assessments. The healthcare organization reserves the right to audit vendors for compliance and to impose penalties for non-compliance.

(4) Data Selection & Acquisition: Create and maintain a data acquisition register. This register should be a structured log that tracks the origin, handling, and intended use of data used in AI systems. Register ensures the data is ethically sourced and used, reliable, and fit for its intended purpose.

(5) Data Quality: Evaluate models for nondiscrimination across race, ethnicity, age, gender, socioeconomic status, etc., to address bias.

(6) Data Quality: Train end users on appropriate model use and known limitations.

(7) Data Quality: Conduct and operationalize recurring quality assessments to ensure AI resources and datasets remain unbiased and anonymized.

(8) Data Quality: Complete anonymization and de-identification validation; HIPAA compliance requires all datasets used for AI training to be de-identified unless explicit patient consent is obtained.

(9) Data Provenance: Maintain oversight and documentation when linking de-identified datasets; assess whether linkage elevates re-identification risk and update HIPAA risk evaluation accordingly.

(10) Data Preparation: Establish a standard operating procedure (SOP) for AI data preparation, including rules for cleaning (e.g., handling missing values), encoding practices (e.g., one-hot, label encoding), normalization or standardization methods, procedures for partitioning (e.g., train, validate, test splits with appropriate temporal or patient separation).

(11) Data Preparation: Establish written policies & procedures for HIPAA-compliant de-identification (Safe Harbor or Expert Determination) while considering state privacy laws.

(12) Data Preparation: Evaluate each PHI-containing dataset and intended use to select the de-identification method that minimizes re-identification risk.

(13) Data Preparation: Maintain continuous oversight of de-identified datasets and re-evaluate risk if datasets are linked or new identifiers emerge.

Subdomain 4.4: Third Party Management

Healthcare delivery organizations should document processes for how roles, responsibilities, and accountability structures are defined and documented between their organization and third parties, when applicable.

Baseline Operational Controls

(1) Include clauses in all AI vendor contracts that define liability for harm or failure resulting from the AI system; outline requirements for disclosure of known model limitations or risks; specify clinical, operational and regulatory responsibilities.

Subdomain 4.5: Feedback and Training

Healthcare delivery organizations should define and document processes for gathering user feedback, incident reporting, and training. Documentation of these processes builds trust, supports safe use, and ensures that relevant stakeholders have the information they need to understand and assess the positive and negative impacts of the AI solution.

Baseline Operational Controls

(1) Ensure all clinical staff receive annual training on AI systems, including their purpose, implementation, and alignment with HITECH Act goals; reinforce this via recurring briefings and optional seminars for deeper engagement.

(2) Provide whistleblower protections for both internal and external reporting, including: anonymity, confidentiality, non-retaliation assurances, and clear disclosure of data retention policies; provide an 'I accept' confirmation at submission to reinforce awareness and compliance.

(3) Ensure data included in incident communication falls under relevant federal laws and protections; ensure awareness of differing state law guidelines.

(4) Notify Head of relevant department(s) after any mandatory reporting incident that occurred; ensure it is known that all parties who cause the incident are made responsible for reporting any voluntary or mandatory incident, otherwise institute is put at risk.

(5) Ensure transparency for the following stakeholders: Patients (via printed materials, patient portals, consent forms, etc.), Clinicians and staff (through embedded resources in clinical tools, EHRs, etc.), Regulators and researchers (through external publication, disclosures, vendor coordination, etc.).

(6) Provide all patients, upon admission, with a comprehensive notice of privacy practices, including how AI systems may use personal health information (PHI), within 30 days in compliance with HIPAA.

(7) Test AI solutions in non-patient environments before full deployment; educate staff on how to use integrated health data platforms with supporting risk assessments and curated training resources, as able.