




CHAI AI Intake Playbook

September 2026

VERSION 1.0 RELEASE

Common Intake Controls & Evidence Framework for Health AI Solutions

A companion to *Subdomain 4.4: Third Party Management Playbook* in the CHAI AI Governance Playbook series. Developed by the CHAI AI Intake Working Group.

Before You Start

For more about the CHAI AI Governance Framework, the methods used to develop the playbook series, and information on how to use the playbooks, please refer to the *CHAI AI Governance Playbooks: Introduction & Background* document. This playbook is a companion to *Subdomain 4.4: Third Party Management Playbook*, which anticipated the need for a standardized AI Intake Questionnaire. This playbook describes the **common AI intake controls** developed by the CHAI AI Intake Working Group: the methodology used to develop them, how institutional questionnaires informed them, how they map to the CHAI Applied Model Card, and the piloting phase that comes next.

This document and the accompanying *CHAI AI Intake Controls and Evidence* workbook are **Version 1.0**, released after review by the AI Intake Working Group. They are expected to evolve with further development, including the pilot phase described in this playbook. It is assumed and appropriate that implementation will differ based on organizational goals, resources, and maturity.

Disclaimer

The 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 input 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 products 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

The *CHAI AI Intake Playbook* presents a standardized set of **57 vendor intake controls across 11 domains**, developed by synthesizing more than 800 intake questions in active use across leading health systems. The controls are designed so that a vendor can gather and document its evidence in a structured response accepted by many potential implementers, while health systems gain a predictable, risk-proportional review structure. This playbook explains how the controls were developed, how they use the CHAI Applied Model Card as the foundational evidence artifact (and where the Model Card leaves gaps), how intake connects to the risk assessment, contracting, and operational control phases described in Subdomain 4.4: Third Party Management Playbook, and what the pilot phase of this work will involve.



Background Context: Why Standardize AI Intake?

AI intake is the first conversation between a healthcare delivery organization (HDO) and an AI vendor, and the first organized look at a solution. Done well, it puts a usable evidence base in front of assessors and directly informs risk assessment, contracting, and operational controls. Done inconsistently, it becomes the single largest source of delay in AI governance.

Today, nearly every HDO maintains its own intake questionnaire. The questions overlap heavily in intent but differ in wording, structure, and evidence expectations. Vendors complete dozens of different questionnaires, and HDO reviewers spend the bulk of their time gathering information and chasing missing documentation rather than

assessing risk. As the volume of vendor AI solutions accelerates, this duplicated effort does not scale for either side of the equation.

Where does intake sit in the broader governance picture? *Subdomain 4.4: Third Party Management Playbook* describes a four-phase workflow that most organizations move through when bringing an AI solution from interest to in-use governance. **This playbook operationalizes Phase 1 to Phase 2.**



The third party management workflow (*Subdomain 4.4: Third Party Management Playbook*). After go-live, monitoring and re-review feeds back into intake whenever there is a material change to the model, the use case, or the regulatory picture.

Note: What Subdomain 4.4: Third Party Management Promised

The Subdomain 4.4: Third Party Management Playbook indicated the need for AI Intake standardization: "the CHAI member community is defining a standard set of risk-based, AI intake questions and controls to streamline this process and ensure that the appropriate documentation and evidence is available for risk determination, solution evaluation, contracting considerations, lifecycle management, and risk mitigation... The primary goal of this standard set of AI intake questions is that vendors would only need to gather and document information once to then be used across intake and procurement processes with HDOs that are aligned to these AI intake questions." This playbook and the accompanying controls workbook deliver the first public version of that standard set.

What We Heard from CHAI Members

To ground the controls in real intake practice, CHAI polled its AI Intake Working Group, with participants including health systems and AI vendors, on how intake works inside their organizations today. The results describe both the challenges and the opportunity for standardization.

The current state of AI intake

- **Median time to approval:** most institutions report between **2 weeks and 2 months** to approve a vendor AI solution.
- **Where review time goes:** respondents report that AI review spends the most time on **information gathering** and **vendor follow-ups for missing information**, not on the assessment itself.
- **Ideal questionnaire size:** most respondents put the ideal intake at roughly **10–40 questions**, far fewer than many questionnaires in circulation today.
- **Prioritization is already happening:** 7 institutions prioritize review of certain solutions over others, expediting solutions that are **low risk, EHR-based, limited in data exposure (minimal or no PHI), high ROI, or mission-critical IT**.

What vendors most often fail to submit

The information most commonly missing from initial vendor submissions maps almost exactly to the evidence expectations in this framework:

- Validation evidence
- Monitoring plan

- Performance metrics
- Subgroup performance
- Workflow integration
- Incident response plan

What reviewers wish they received

Asked what information they do *not* receive today that would make it easier to expedite vendors through governance, members named:

- Clearly defined scope boundaries (what the model does vs does not do)
- Hallucination rate (for generative AI)
- Subgroup performance metrics with numeric thresholds
- Clear disablement ("off switch") capabilities

The top information priorities during intake

Members ranked the five most important categories of information received during the intake process:

1. Clinical safety and effectiveness
2. Legal and regulatory compliance
3. Data integrity & lineage
4. Model validation & monitoring
5. Fairness & non-discrimination

Note: The Mandate for a Common Application

Members overwhelmingly cite the rapid growth of vendor AI solutions as the force straining today's intake processes, and the shared resources they say would help most are precisely the components of this framework: a common vendor intake questionnaire, a control framework and review rubric, risk tiering and triage logic, an evidence checklist, and shared templates such as a standardized model card. A majority of members polled said they would trust a "fast pass" vendor application completed once and reused across institutions, provided it is standardized and complete.

Methodology: How We Got to the Control Statements

The common intake controls are the product of a structured synthesis of real institutional questionnaires followed by iterative expert review. The Working Group is composed of clinical, technical, governance, legal, and procurement subject-matter experts drawn from CHAI member organizations, including health systems, AI vendors, payers, and academic institutions. The process moved through five steps.

Step 1 – Gather data

The working group collected **more than 800 individual intake questions from approximately 16 questionnaires** in active use across participating health system members, supplemented by inputs from the CHAI Risk Workgroup. Contributing systems included UPenn (Penn Medicine), UNC Health, Mass General Brigham, Mayo Clinic, Stanford Health Care, Tampa General Hospital, Mercy, MD Anderson, Ascension, AdventHealth, and Inova. After deduplication, **761 source questions** formed the evidence base.

Step 2 – Cluster and categorize

Source questions were clustered into approximately 260 concept clusters and organized into **11 control domains**. Initial machine-assisted clustering was audited in full by the working group; that review identified and corrected systematic errors, including over-assignment of questions to a governance catch-all category, over-merged clusters combining distinct concepts, and unevaluated alternative question phrasings. Every cluster assignment used in the final framework reflects human subject-matter review.

Step 3 – Draft controls

Each cluster of questions was converted from interrogative form into a **plain-language control statement designed for a yes/no vendor response** (e.g., "Intended use, patient population, and care setting are documented."). Instead of every institution asking its own open-ended variant of the same question, the vendor attests to a shared statement and points to the evidence behind it. Controls that surface only for high-risk deployments are explicitly flagged.

Step 4 – Iterate

The draft controls went through iterative review across multiple working group sessions, each focused on a subset of the controls. Feedback was incorporated through revised statements, consolidation of overlapping controls, attestation bundling proposals, and the addition or retirement of individual controls (the workbook preserves a full mapping from each control to its previous identifier). Controls will continue to be updated and iterated upon through the pilot phase.

Step 5 – Architect the evidence

Finally, each control was mapped to the evidence that satisfies it: an evidence type, illustrative artifact examples, an attestation ruling (whether vendor signature alone closes the control), a mapping to the CHAI Applied Model Card, and a mapping to the risk-categorization questions the control helps answer. Related artifacts were then bundled into 13 evidence packages so that one submission satisfies many controls.

Design Principles

Four principles governed every drafting and consolidation decision:

- **Risk-proportional application.** Controls scale in evidence depth with the AI solution's risk tier. Low-risk deployments encounter only the baseline set; Medium and High deployments encounter progressively more rigorous evidence expectations.
- **Evidence standardization.** Controls are written to be answered yes/no and satisfied by defined artifact types, so vendors can anticipate exactly what to prepare and reviewers can anticipate exactly what they will receive.
- **Evidence reuse.** Controls are mapped to the CHAI Applied Model Card wherever possible so vendors are not asked to produce overlapping artifacts. The Model Card serves as the foundational artifact for the intake review.
- **Maximize attestation.** Where risk permits, a vendor signature substitutes for individual artifact review. Attestation bundles collapse multiple low-risk control responses into a single vendor signature ceremony, reducing operational burden while preserving control fidelity.

The Framework at a Glance

The result is a set of **57 controls across 11 domains**, published in the companion *CHAI AI Intake Controls and Evidence* workbook. Vendors complete the controls; implementers apply them in proportion to their own risk assessment of the solution.

16

baseline controls asked of every vendor, regardless of tier

41

mitigation controls triggered by the deployment's risk profile

28 / 47 / 57

controls applied at the Low / Medium / High risk tiers, respectively

40%

of controls closed by vendor attestation alone, with no artifact review

Eleven control domains

#	Domain	Prefix	Example control
1	Clinical Safety & Effectiveness	CSE	Known risks, limitations, and out-of-scope settings are disclosed.
2	Data Integrity & Lineage	DIL	Training data sources are documented.
3	Fairness & Non-Discrimination	FND	The AI system has been tested for performance differences across demographic and clinical subgroups.
4	Governance Structure & Oversight	GSO	An escalation path is documented for issues or concerns about the tool.
5	Impact Measurement & ROI	IMR	Outcome data is collected and reviewed on a defined schedule.
6	Legal & Regulatory Compliance	LRC	Vendor discloses how customer data is used, including any uses for product improvement or model training.
7	Model Validation & Performance Monitoring	MVP	The vendor validates model performance on a representative dataset for the customer's intended use prior to deployment.
8	Security & Access Control	SAC	The vendor holds, or is actively pursuing, recognized security certifications appropriate to the data handled.
9	System Reliability & Resilience	SRR	Failover, downtime, and AI-service contingency procedures are documented.
10	Transparency & Traceability	TRT	Third-party components and foundation models are disclosed in sufficient detail to support security and BAA assessment.
11	User Experience & Workflow Integration	UWI	The vendor documents how the tool integrates with required customer systems and workflows.

Risk-tiered application

Each control carries the minimum risk tier at which it applies, using tier definitions aligned with CHAI's Risk Categorization Tool (v3) (additional detail in *Subdomain 4.2: Risk and Impact Assessments Playbook*). In brief: **Low**-risk tools present or process information with no autonomous action and full human review of outputs;

Medium-risk tools influence clinical or operational workflows in ways that may be acted on without independent verification or handle sensitive patient data; **High**-risk tools directly support diagnosis, treatment, or triage decisions, operate autonomously or semi-autonomously, or are deployed in high-acuity settings. Tier assignment is determined by the health system at intake, with vendor input where relevant, and clinical safety controls are excluded for non-clinical tools. High-tier control statements are visibly flagged ("HIGH") so vendors can see at a glance which expectations attach only to high-risk deployments.

The attestation ladder

Every control carries one of three evidence rulings, applied in proportion to risk:

- **Attestation sufficient (23 controls).** Vendor signature alone closes the control; no separate evidence review is required. Attestation bundles collapse multiple such responses into a single signature ceremony.
- **Attestation + supporting document (17 controls).** Signature plus a referenced artifact, such as a Service Level Agreement (SLA) excerpt or workflow diagram.
- **Evidence required (17 controls).** A substantive artifact is required; attestation alone is not sufficient. Concentrated in medium- and high-risk mitigation controls such as validation studies, bias evaluations, and failure-mode analyses.

Thirteen evidence packages

Rather than requesting artifacts control-by-control, related evidence is bundled into packages that mirror how vendors actually maintain documentation. Six baseline packages are expected of all vendors; the remainder are triggered by risk tier.

Package	Trigger	Name	Example artifacts
B1	Baseline – all vendors	CHAI Applied Model Card	Completed CHAI Applied Model Card
B2	Baseline – all vendors	Business Associate Agreement (BAA) / Privacy & Data Use Agreement (DUA)	BAA template, privacy policy, DUA
B3	Baseline – all vendors	Security Certification Pack	SOC 2 / HITRUST / ISO 27001, audit report, certification roadmap
B4	Baseline – all vendors	Operations & Monitoring Documentation	SLAs, lifecycle policy, monitoring plan, governance documentation
B5	Baseline – all vendors	Workflow & Integration Documentation	Workflow diagrams, integration and architecture documentation
B6	Baseline – all vendors	Validation Plan / Proof of Concept (POC) Scope	Validation plan, proof-of-concept scope, benchmark results
L1	Low tier	Acceptable Use & Training Pack	Acceptable use policy, training materials, escalation protocol
M1	Medium tier	Validation & Bias Package	Validation study, subgroup testing results, bias methodology
M2	Medium tier	Operational Resilience Package	Continuity plan, dependency inventory, incident response plan
M3	Medium tier	Usability Evidence	Usability study, override/dismissal analytics, training documentation
H1	High tier	Safety & Failure Analysis Package	Safety surveillance plan, Failure Mode and Effects Analysis (FMEA), human factors study

H2	High tier	AI Security Assessment Package	AI security assessment, red-team summary, threat model
H3	High tier	Transparency & Control Package	Audit log specification, confidence-score UX, suspension policy

Note: Evidence Examples Are Illustrative, Not Exhaustive

Artifact examples in the workbook illustrate what typically satisfies each control; vendors may submit equivalent artifacts that map to the same intent. For example, a SOC 2 Type II report, a HITRUST CSF certification, or an active-pursuit certification roadmap all satisfy the security-certification expectation.

The 57 Common Intake Controls

The full control set is listed below by domain, exactly as published in the companion workbook. Controls marked with an asterisk (*) are **baseline controls** asked of every vendor and required for risk assessment; all others are mitigation controls triggered by assessed risk level. **Tier** is the minimum risk level at which the control applies (High-tier controls surface only for high-risk deployments; clinical safety controls do not apply to non-clinical tools). **Evidence ruling** indicates whether vendor signature alone closes the control ("Attestation"), signature plus a referenced artifact is expected ("Attest. + doc"), or a substantive artifact is required ("Evidence"). **Model Card** shows the degree to which the CHAI Applied Model Card already carries the information needed to satisfy the control. Conditional Security & Access Control (SAC) controls apply only when SAC-01 is answered "No."

Clinical Safety & Effectiveness

ID	Tier	Control statement	Evidence ruling	Model Card
CSE-01 *	Low	Intended use, patient population, and care setting are documented.	Attestation	Full
CSE-02 *	Low	Known risks, limitations, and out-of-scope settings are disclosed.	Attestation	Full
CSE-04 *	Medium	A bi-directional process exists for the customer and vendor to identify, track, and report adverse events or safety/quality concerns.	Attest. + doc	Partial
CSE-07 *	Medium	Human oversight processes are incorporated into the tool's workflow and documented.	Attest. + doc	Partial
NEW-02 *	Low	The time available for clinician review of AI output before a clinical decision or action is required is characterized.	Attestation	Partial
CSE-03	Medium	Tool performance has been validated for the intended or a similar care or operational setting, with key metrics documented.	Evidence	Partial
CSE-05	High	A continuous safety surveillance and post-market monitoring plan is in place.	Evidence	Partial
CSE-06	High	Failure modes and mitigations are documented for the tool.	Evidence	Partial
CSE-08	Medium	User override and dismissal rates for AI outputs are tracked and used to support appropriate trust calibration.	Attest. + doc	Partial
CSE-09	High	The tool is demonstrated to support user decision-making, with evidence from evaluation in high-risk settings.	Evidence	Partial
CSE-10	Low	The vendor documents how the tool is intended to integrate into customer workflows across applicable user roles.	Attest. + doc	Partial

Data Integrity & Lineage

ID	Tier	Control statement	Evidence ruling	Model Card
DIL-01 *	Low	Training data sources are documented.	Attestation	Full
DIL-02 *	Low	Data quality and preprocessing steps are documented.	Attestation	Full
DIL-03	Low	Data inputs required at the point of use and associated outputs are clearly specified.	Attestation	Full
DIL-04	High	Data quality gating conditions for safe operation are defined, and quality status is communicated to users.	Evidence	Partial

Fairness & Non-Discrimination

ID	Tier	Control statement	Evidence ruling	Model Card
FND-02 *	Medium	The vendor identifies biases observed during development and documents mitigation steps taken.	Attest. + doc	Full
FND-01	Medium	The AI System has been tested for performance differences across applicable demographic and clinical subgroups, with methodology and results documented. This includes known biases and limitations.	Evidence	Partial
FND-03	High	Vendor monitors performance post-deployment and shares results with the customer per agreed cadence and contractual terms.	Evidence	Partial

Governance Structure & Oversight

ID	Tier	Control statement	Evidence ruling	Model Card
GSO-02 *	Low	A lifecycle management process exists for the tool covering adoption, versioning, and retirement.	Attest. + doc	None
GSO-01	Low	A designated accountable party at the vendor is assigned to oversee the tool's deployment, performance, and compliance.	Attestation	None
GSO-03	Low	Acceptable use guidance is documented and provided to customer organizations for distribution to users.	Attest. + doc	None
GSO-04	Low	An escalation path is documented for issues or concerns about the tool.	Attest. + doc	None
GSO-05	Medium	The vendor supports a defined re-review cycle and provides updated documentation on a regular schedule.	Attestation	Partial
GSO-06	High	Conditions that trigger tool suspension or removal are defined.	Evidence	Partial
GSO-07	High	Override, escalation, and rapid-shutdown mechanisms are documented.	Evidence	None

Impact Measurement & ROI

ID	Tier	Control statement	Evidence ruling	Model Card
IMR-01	Low	The vendor will provide outcome metrics agreed upon with the customer prior to deployment.	Attestation	Partial
IMR-02	Medium	Outcome data is collected and reviewed on a defined schedule.	Attest. + doc	Partial
IMR-03	Medium	The vendor provides data, dashboards, or APIs to support outcome measurement at a cadence agreed with the customer.	Attest. + doc	Partial

Legal & Regulatory Compliance

ID	Tier	Control statement	Evidence ruling	Model Card
LRC-02 *	Low	Vendor demonstrates and commits to HIPAA-compliant data handling, through a Business Associate Agreement (BAA) for example. This includes prohibiting re-identification of de-identified data.	Attestation	Partial
LRC-01	Low	Applicable regulatory clearances or approvals are disclosed, with current status. (N/A permitted for non-clinical use cases.)	Attestation	Full
LRC-03	Medium	Vendor actively monitors for bias, drift, and takes pre-defined and appropriate action to mitigate.	Evidence	None
LRC-04	Low	Vendor discloses how customer data is used, including any uses for product improvement or model training.	Attestation	None
LRC-05	Low	Conflicts of interest related to the tool are disclosed.	Attestation	Partial

Model Validation & Performance Monitoring

ID	Tier	Control statement	Evidence ruling	Model Card
MVP-01 *	Low	The vendor validates model performance on a representative dataset for the customer's intended use prior to deployment.	Attestation	Partial
MVP-03 *	Medium	The vendor confirms post-deployment performance monitoring capability; specific metrics and thresholds are specified with the customer.	Attest. + doc	Partial
MVP-02	Medium	External, prospective, or similarly-situated-site validation has been conducted where appropriate.	Evidence	Partial
MVP-04	Medium	A model drift or degradation response process is defined.	Attest. + doc	Partial
MVP-05	Low	Version history and change records are maintained and shared with the customer.	Attestation	Full

Security & Access Control

ID	Tier	Control statement	Evidence ruling	Model Card
NEW-04 *	Low	The model's training and update approach post-deployment is disclosed, including whether the model is static, periodically retrained, or continuously learning from live production data.	Attestation	Partial
SAC-01 *	Low	The vendor holds, or is actively pursuing, recognized security certifications appropriate to the data handled.	Evidence	Full
SAC-02	Low	[CONDITIONAL - if SAC-01 response is No] Access to the tool and its outputs is role-based.	Attestation	None
SAC-03	Low	[CONDITIONAL - if SAC-01 response is No] Data is encrypted in transit and at rest.	Attestation	None
SAC-04	Medium	[CONDITIONAL - if SAC-01 response is No] A security incident response process is in place.	Evidence	None
SAC-05	High	AI-specific security risks are documented and mitigated, for example prompt injection, jailbreaking, data exfiltration via outputs, and third-party LLM dependencies.	Evidence	None

System Reliability & Resilience

ID	Tier	Control statement		Model Card
----	------	-------------------	--	------------

			Evidence ruling	
SRR-01	Medium	Failover, downtime, and AI-service contingency procedures are documented.	Attest. + doc	None
SRR-02	Medium	System performance and incident communications are reported to the customer on a defined basis.	Attestation	None
SRR-03	Medium	Third-party AI service dependencies and contingency plans are documented.	Attest. + doc	None

Transparency & Traceability


ID	Tier	Control statement	Evidence ruling	Model Card
TRT-05 *	Medium	Tool outputs can be explained at a level appropriate for the user's role and the decision being supported.	Attest. + doc	Partial
TRT-02	Low	The model type and key development decisions are disclosed.	Attestation	Full
TRT-03	Low	Third-party components and foundation models are disclosed in sufficient detail to support security and BAA assessment.	Attestation	Full
TRT-04	Low	Material version updates and any change to underlying foundation models are communicated to the customer in advance.	Attestation	Full
TRT-06	High	Uncertainty or confidence scores are displayed to users alongside outputs in human-in-loop scenarios.	Evidence	Partial
TRT-07	High	For agentic systems, confidence scores and decisions are logged in an audit trail with periodic monitoring.	Evidence	Partial

User Experience & Workflow Integration

ID	Tier	Control statement	Evidence ruling	Model Card
UWI-02 *	Low	The vendor documents how the tool integrates with required customer systems (e.g., EHR, identity provider, data feeds) and provides evidence of integration approach.	Attest. + doc	None
UWI-01	Low	The vendor provides training materials and curriculum suitable for the customer organization to deliver to end users.	Attest. + doc	Partial
UWI-03	Low	The vendor documents both customer-service feedback channels and AI-output feedback mechanisms.	Attestation	None
UWI-04	Medium	Usability evidence is provided commensurate with the vendor's maturity.	Evidence	Partial

The CHAI Applied Model Card: Foundation and Gaps

The single most important evidence-reuse decision in this framework is anchoring intake on the **CHAI Applied Model Card**. Every control was assessed for the degree to which a completed Model Card already carries the information needed to satisfy it, and where coverage exists, the workbook names the specific Model Card field(s) a reviewer should consult. The result:



In other words, **72% of the common intake controls are at least partially satisfied by a single, standardized artifact the vendor completes once.** A vendor that submits a complete CHAI Applied Model Card closes 12 controls outright and materially advances 29 more. The 12 include intended use and population, known risks and limitations, training data sources, model type, foundation-model and third-party component disclosure, version history, regulatory status, and security accreditations.

What the Model Card covers well

The Applied Model Card is strongest exactly where intake reviewers rank their information priorities highest: descriptive transparency about the solution. Intended use and workflow, targeted patient population, known risks and cautioned out-of-scope settings, development data characterization, bias mitigation during development, model type and key development decisions, foundation models and third-party components, release stage and version history, and security/compliance accreditations are all first-class Model Card content.

Where the Model Card does not currently reach

Sixteen controls have **no current Applied Model Card coverage**. They concentrate in four areas. Notably, all are operational rather than descriptive: they concern how the vendor will behave over the life of the relationship, not what the model is.

- **Operational governance (GSO).** A designated accountable party at the vendor; a lifecycle management process covering adoption, versioning, and retirement; acceptable use guidance; a documented escalation path; and, for high-risk deployments, override, escalation, and rapid-shutdown ("off switch") mechanisms.
- **Security operations (SAC).** Role-based access control, encryption in transit and at rest, and a security incident response process (where recognized certifications are absent), plus documentation and mitigation of AI-specific security risks such as prompt injection, jailbreaking, and data exfiltration.
- **System reliability & resilience (SRR).** Failover, downtime, and AI-service contingency procedures; performance and incident reporting to the customer; and third-party AI service dependencies with contingency plans.
- **Integration, data use & feedback (UWI / LRC).** Integration documentation for required customer systems, customer-facing feedback channels, disclosure of how customer data is used (including for product improvement or model training), and active bias/drift monitoring commitments.

Partial coverage: where depth is the gap

For the 29 partially covered controls, the Applied Model Card typically names the topic but does not carry the depth an HDO needs. Ongoing maintenance fields gesture at monitoring, but not at specific metrics, thresholds, reporting cadence, or drift-response processes. Fairness fields describe bias mitigation during development, but not subgroup performance metrics with numeric thresholds. Intended-use fields describe the workflow, but not the time available for clinician review before a decision is required, nor human oversight process documentation.

Note: The Gaps Are Exactly What Reviewers Say Is Missing

The working group poll asked reviewers what they do not receive today that would expedite governance review. Their answers sit almost entirely in the Model Card's partial- and no-coverage zones: clearly defined scope boundaries, hallucination rates for generative AI, subgroup performance metrics with numeric thresholds, and clear disablement ("off switch") capabilities. The same is true of the items most commonly missing from initial submissions: validation evidence, monitoring plans, performance metrics, subgroup performance, workflow integration, and incident response plans. The coverage mapping and the field poll arrived at the same gaps independently. That convergence is strong evidence the intake controls target the right gaps, and it gives the Applied Model Card a concrete agenda for future versions.

The framework's posture toward these gaps is pragmatic: the Model Card remains the foundational artifact, the 16 uncovered controls are satisfied through lightweight attestations and artifacts vendors already maintain (SLAs, runbooks, incident response plans, integration guides), and the coverage mapping gives the Model Card's stewards a prioritized list of candidate fields for future versions.

From Intake to Contract: Powering the Rest of Subdomain 4.4: Third Party Management

Standardized intake is not an end in itself. Its value is measured by how much it improves the three phases that follow it in the Subdomain 4.4: Third Party Management Playbook workflow. Each control in the workbook is explicitly mapped to the downstream decisions it informs.

Phase 2 – Risk assessment

Each baseline control is mapped to the risk categorization question(s) it answers (e.g., distance from patient, decision autonomy, consequences of failure or error, data transparency, monitoring difficulty, use of sensitive data). This means a completed intake is not merely a document collection exercise: it directly populates the risk categorization described in *Subdomain 4.2: Risk and Impact Assessments Playbook*, which in turn determines which mitigation controls and evidence packages apply. Intake and risk assessment reinforce each other: the risk tier trims the control set, and the control responses substantiate the tier.

Phase 3 – Contracting

Subdomain 4.4: Third Party Management Playbook is explicit that AI addenda are starting points whose substance is shaped by intake and risk assessment. The intake controls generate contracting inputs directly:

- **Disclosure obligations.** Controls on known limitations, training data, and foundation-model dependencies (CSE-02, DIL-01, TRT-03) establish the disclosure baseline that the contract then makes ongoing, including notice of material model changes and regulatory reclassification.
- **Monitoring and SLA terms.** Controls confirming post-deployment monitoring capability and drift response (MVP-03, MVP-04, FND-03) surface exactly which metrics, thresholds, reporting cadences, and remedies the SLA should specify.
- **Incident and adverse-event terms.** The bi-directional adverse event process (CSE-04) and incident reporting controls give legal teams the factual basis for allocating reporting duties and timelines.
- **Data use terms.** The customer data use disclosure (LRC-04) and BAA controls (LRC-02) feed the data use restrictions and AI-specific BAA addendum content described in Subdomain 4.4.
- **Gaps become clauses.** Where a vendor answers "no" or submits thin evidence, that gap converts into a contract term: a cure obligation, a supplemental representation, a conditional approval with a defined remediation timeline, or a walk-away trigger. Intake findings demonstrably shape the final contract, whether they precede signature or drive a subsequent amendment.

Phase 4 – Operational controls

Risks that survive contracting are managed at the workflow level, and intake identifies where those safeguards are needed. Controls on human oversight processes (CSE-07), clinician review time (NEW-02), user training (UWI-01), override and dismissal tracking (CSE-08), and rapid-shutdown mechanisms (GSO-07) tell the deploying organization which operational control patterns the deployment requires: human review of outputs, default-off postures, named owners and escalation paths, and independent rollback procedures.

Worked Illustration: One Intake, Three Downstream Uses

Scenario: A health system receives a completed common intake from a vendor of an ambient documentation tool. The Applied Model Card (package B1) closes the descriptive controls. The intake reveals two findings: the vendor's monitoring plan names metrics but no thresholds (MVP-03, partial), and no rapid-shutdown mechanism is documented (GSO-07, absent).

Downstream: The risk assessment classifies the tool Medium risk, activating the M1–M3 evidence packages. Contracting converts the monitoring gap into an SLA term with defined thresholds, reporting cadence, and service-credit remedies. The shutdown gap becomes an operational control: the deployment goes live with an independent, HDO-owned disablement procedure and a named escalation owner, with a contractual commitment from the vendor to deliver a documented disablement capability in its next release. Every downstream decision traces to a specific intake control.

Next Phase: Piloting the Framework

Version 1.0 is a starting point, and CHAI intends to keep refining it. The next phase of the AI Intake Working Group's effort is a structured pilot: putting the common controls, evidence packages, and Applied Model Card mappings in front of real vendor–implementer pairs and measuring whether they deliver the burden reduction they promise.

What the pilots are designed to test

- **Vendor burden.** How long does it take a vendor's technical team to complete the controls once, including assembling the evidence packages? Which controls generate confusion or disproportionate effort?
- **Evidence sufficiency.** Do the artifact examples and attestation rulings give reviewers what they actually need, or do they still generate follow-up requests? Every follow-up in a pilot is a signal that a control statement or evidence expectation needs revision.
- **Cycle time.** Does a standardized submission move a solution through governance measurably faster than the 2-week-to-2-month baseline reported by the working group?
- **Attestation tolerance.** Are implementers comfortable closing the designated low-risk controls on signature alone? Where does attestation prove insufficient in practice?
- **Applied Model Card coverage in practice.** Does the Full/Partial/None mapping hold up against real completed Model Cards, and which partial-coverage fields most need deepening?
- **Reusability across institutions.** The central promise: can a single vendor submission be accepted by multiple implementers with different governance structures, risk tolerances, and maturity levels?

How the pilot feeds revision

Pilot findings will be incorporated through the same revision discipline used to develop the controls: revised statements, consolidation candidates, attestation bundling adjustments, and the addition or retirement of individual controls, with identifier continuity preserved in the workbook. The coverage mapping will also be updated against the current version of the Applied Model Card, and gap findings will be shared with the Applied Model Card's stewards as candidate fields for future versions.

Note: How to Participate

CHAI is seeking healthcare delivery organizations and AI vendors to participate in the pilot phase, particularly pairs already engaged in an active procurement and organizations of varying size and governance maturity, so the framework is stress-tested beyond large academic medical centers. Working Group feedback on individual controls, evidence expectations, or consolidation candidates should be directed to the AI Intake Working Group and will be incorporated into the next revision cycle. Contact: greg@chai.org.

Implementation Guiding Questions

For healthcare delivery organizations

- How closely does your current intake questionnaire map to the 11 domains and 57 controls? Which of your questions have no counterpart here, and is that a local necessity or an artifact of history?
- Could you accept a vendor's standardized submission, completed for another institution, as the starting evidence base for your own review? If not, what specifically would you need added?
- Where does your review time actually go today: assessment, or information gathering and vendor follow-up? What would change if initial submissions arrived complete?
- Which controls would your organization be comfortable closing on vendor attestation alone, and at which risk tiers?
- Do your intake findings today demonstrably shape contract terms, as Subdomain 4.4 recommends? Which intake gaps have you converted into clauses, and which were lost between intake and signature?
- Do you require the CHAI Applied Model Card (or equivalent) as a condition of contract execution today?

For AI vendors

- Could your technical team complete the 57 controls today? Which evidence packages (B1–B6) could you assemble this quarter, and which would require net-new documentation?
- Is your CHAI Applied Model Card complete and current, and do you have supplemental artifacts ready for the operational areas the Model Card does not cover: lifecycle management, escalation paths, resilience procedures, and AI-specific security assessment?
- Can you state your solution's scope boundaries, subgroup performance with numeric thresholds, and (for generative AI) hallucination or error rates? These are the items reviewers most often request and least often receive.
- Who is your designated accountable party for a customer deployment, and is that documented somewhere a customer can reference?

Tools & Resources

The following resources are available on the CHAI website under the AI Intake work group resources.

CHAI AI Intake Controls and Evidence Workbook

- The companion workbook to this playbook: 57 controls with control statements, risk tiers, evidence packages and examples, attestation rulings, Applied Model Card coverage and field mappings, and risk-question mappings. Version 1.0, 2026.

CHAI Applied Model Card (AMC)

- A standardized documentation framework that describes an AI solution in the context of a specific health use case. The foundational evidence artifact (package B1) for the intake review. Available via the CHAI Registry.

Subdomain 4.4: Third Party Management Playbook

- Guidance on the four-phase third party management workflow, AI addenda and contract terms, vendor negotiation preparation, post-deployment enforcement, and operational control patterns. Part of the CHAI AI Governance Playbook series.

Subdomain 4.2: Risk and Impact Assessments Playbook

- Guidance and structure for risk categorization and risk assessment, which determine the tier at which the intake controls apply.

CHAI AI Governance Playbooks: Introduction & Background

- The recommended starting point for the playbook series, covering the CHAI AI Governance Framework and the methods used to develop the playbooks.

Attribution

This playbook was developed by the CHAI AI Intake Working Group, composed of clinical, technical, governance, legal, and procurement subject-matter experts drawn from CHAI member organizations, including health systems, AI vendors, payers, and academic institutions. CHAI extends its sincere appreciation to the health systems that contributed their institutional intake questionnaires, the raw material for the common controls, and to all working group participants for their time, expertise, and thoughtful feedback across the iterative review sessions.

Contributing Organizations

Individuals from the following organizations participated in the AI Intake Working Group. Listing reflects participation only and does not imply endorsement of this playbook by any organization.

AdventHealth	Kaiser Permanente	Oracle
AgileMD	Lumeris	Penn Medicine
AlgoHealth	Marisa	Signall
Ascension	Mass General Brigham	Solventum
Atropos Health	Mayo Clinic	Stanford Health Care
Bayesian Health	MCG Health	Suki
Citadel AI	MD Anderson Cancer Center	Tampa General Hospital
Credo AI	Medasource	UNC Health
Folks	Mercy	UnitedHealth Group
Inova	OnBoard AI	Yale New Haven Health
Iowa Primary Care Association	Optum	

License

This document is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 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. **NoDerivatives:** 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/.

Copyright 2026 © Coalition for Health AI, Inc.