Subdomain 4.4: Third Party Management Playbook

Guiding responsible AI adoption

Before You Start

For more about the CHAI AI Governance Framework, methods used to develop these playbooks, and information on how to use them, please refer to the CHAI AI Governance Playbooks: Introduction & Background document. Note that **baseline operational controls** were consensus-defined by Healthcare Delivery Organizations (HDOs) as "must-haves" for health AI governance. The **implementation guidance** that accompanies the baseline operational controls are **suggestions/recommendations** based on known practices. It is assumed and appropriate that implementation will differ based on organizational goals, resources, and maturity.

Disclaimer

The implementation guidance, frameworks, and examples contained herein represent generalized practices and considerations drawn from across the field of responsible AI deployment in healthcare. They are offered as one set of perspectives to inform your thinking not as the "right way" or "only way" to build, govern, or operate an AI program.

Every organization operates within its own unique context: differing patient populations, clinical priorities, regulatory environments, resource constraints, technical infrastructure, cultural norms, and stage of AI maturity. While the specific methods, timelines, and depth of implementation will appropriately vary across organizations, the underlying controls, principles, risks, and considerations raised throughout this playbook are intended to be engaged with thoughtfully and in good faith by any organization deploying AI in healthcare.

Organizations should prioritize state and federal regulatory requirements, especially for AI solutions that meet the definition of a Software as a Medical Device (SaMD), as defined by the FDA.

Not a Standard of Care. This playbook does not establish, define, or reflect a legal, regulatory, professional, or industry standard of care, nor is it intended to serve as one. It is a voluntary, aspirational, and educational resource designed to support organizational learning and dialogue. The field of healthcare AI is rapidly evolving, and reasonable organizations and professionals may differ on how to approach the issues discussed herein. The practices described may not be appropriate, feasible, or advisable for every organization, every use case, or every point in time. Adoption of any particular practice described in this playbook is entirely voluntary, and a decision to adopt, adapt, partially implement, or decline to implement any practice should not be construed as evidence of compliance with, or deviation from, any standard of care, duty, or legal obligation. This playbook is not intended to be used, and should not be used, as a benchmark, checklist, or measuring stick against which the conduct of any

organization, clinician, or other party is evaluated in any litigation, regulatory proceeding, accreditation review, or similar context.

General Informational Purpose. The information in this playbook is provided for general informational and internal operational purposes only. It is intended to support organizations in thinking through the design and stewardship of their AI programs and should be used in conjunction with from clinical leaders, technical experts, compliance officers, ethicists, patient and community representatives, and other advisors familiar with the specifics of your organization.

Not Legal Advice. Nothing in this playbook constitutes legal advice, and it should not be relied upon as a substitute for advice from a qualified attorney licensed to practice in the relevant jurisdiction. Nothing in this document creates an attorney-client relationship between the reader and CHAI, its officers, directors, employees, or contributors. The contents reflect general information and observed practices as of the date of publication and may not account for changes in law, regulation, technology, or circumstance that occur thereafter.

No Warranty. While reasonable efforts have been made to ensure accuracy, no representation or warranty is made as to the completeness, timeliness, accuracy, or suitability of the information for any particular purpose. Laws, regulations, and best practices vary by jurisdiction and evolve over time, and the application of those laws and practices to specific facts can lead to different outcomes.

Illustrative Examples. Any tools, resources, examples, or technology product described or presented are for educational and illustrative purposes only and do not represent endorsement by CHAI or guarantee specific results. Readers should consult qualified legal counsel and relevant experts for specific terms, products, or services.

Independent Judgment Required. Readers should consult with qualified legal counsel, clinical experts, and other appropriate professional advisors before taking, or refraining from taking, any action based on the information in this playbook. The decision to adopt, modify, or disregard any practice described herein rests solely with the reader and their organization. CHAI expressly disclaims any and all liability for any loss, damage, or other consequence, including but not limited to claims of negligence, breach of duty, or failure to meet any standard, arising directly or indirectly from reliance on, use of, citation to, or reference to the contents of this document.

At a Glance

Subdomain 4.4: Third Party Management Playbook is designed to guide HDOs to document processes for how roles, responsibilities, and accountability structures are defined and documented between the HDO and third party vendors. This playbook offers actionable implementation steps, real-world examples, and tools to navigate modern AI governance.

Background Context for Subdomain 4.4: Third Party Management

The increasing integration of artificial intelligence into healthcare delivery has changed the nature of vendor relationships. AI solutions enter HDOs through different ingestion pathways, two of which are applicable in the

context of this playbook: (1) vendor or co-developed models/solutions and (2) AI-enhanced features in existing health IT systems. Subdomain 4.4: Third Party Management of CHAI's AI Governance Framework addresses how HDOs should document and manage the roles, responsibilities, and accountability structures that govern relationships with third party AI vendors and partners.

The second pathway, AI-enhanced features in existing health IT systems, is likely responsible for a rapidly growing share of the AI footprint at most HDOs, and organizations often have limited leverage over these embedded capabilities, particularly when they are activated as part of routine platform updates. For this pathway, the primary contractual objective shifts from negotiating bespoke clauses toward two complementary goals: (1) obtaining the vendor's commitment to continually inform the customer of new features, artifacts, and significant updates before they become active; and (2) ensuring that acceptable blanket standard operating procedures (SOPs) govern any solution introduced through the platform. Risk categorization of embedded solutions (see Subdomain 4.2: Risk & Impact Assessments Playbook) is especially valuable in this pathway because it trims the scope of deeper oversight to a more manageable volume and helps prioritize attention on the features that most warrant it.

To fully appreciate the importance of Subdomain 4.4: Third Party Management, it should be understood in the context of Subdomain 4.1: Responsible AI Lifecycle Management and Use Playbook. Subdomain 4.1 establishes that AI solutions should be governed across a defined lifecycle for all relevant ingestion modalities. The AI lifecycle begins with intake and procurement, and this is precisely where Third Party Management has significant impact. Contractual protections established at intake set the foundation for how the HDO will be able to govern AI across the entire lifecycle. Contractual controls serve as enforceable mechanisms across the AI lifecycle.

Note: Embedded AI and Enterprise SaaS at Scale

Many AI features arrive into HDOs through routine software updates from incumbent vendors (e.g., EHR, imaging, RCM, patient engagement platforms). These features frequently bypass traditional procurement gates because no new contract, statement of work, or BAA is executed at the moment the feature is enabled. HDOs should consider:

- **Inventory AI features in existing enterprise platforms** on a recurring cadence, with explicit owners assigned per platform.
- **Re-read each platform's data processing terms specifically for AI model training and derivative use provisions**, which often differ materially from the general terms of service and may be updated independently of the underlying agreement.
- **Apply retroactive governance review** to legacy AI features that were enabled before an AI governance program existed.
- **Define a "default-off" posture** for newly released AI features on enterprise platforms wherever the platform permits, pending governance review.

Retroactive Review Workflow *(example, illustrative-only)*

- **Step 1 - Discover:** Recurring sweep across the enterprise platform list. Use vendor admin consoles and release notes.
- **Step 2 - Triage:** Categorize each feature by risk (see Subdomain 4.2: Risk and Impact Assessment playbook). Default-off for any feature pending review where the platform permits.
- **Step 3 - Re-read terms:** Specifically the AI-related data processing addendum, model training provisions, and other relevant data clauses. These often differ from the general terms of service and may have been updated without notice.
- **Step 4 - Assign owner:** Every embedded AI feature gets a named org owner.
- **Step 5 - Document:** Add to AI Vendor Register with a recurring re-review cadence.
- **Step 6 - Invoke contract levers:** Use change-notification clauses on the contract agreement; request written confirmation that organization content is not used to train shared models (as needed); request deletion / opt-out (where available).

Guidance on Existing Infrastructure

CHAI does not recommend that HDOs restructure existing vendor management, legal, compliance, or procurement processes. Instead, this playbook provides guidance on amending or supplementing existing frameworks to incorporate AI-specific components. This may take the form of:

- Adding an AI Appendix or AI Addendum to new or existing Business Associate Agreements (BAAs) with AI vendors.
- Incorporating AI-specific intake questions into existing procurement workflows.
- Updating standard vendor contract templates to include disclosure and responsibility clauses.
- Assigning AI governance oversight responsibilities within existing compliance, legal, or clinical informatics functions.

Note: Third Party Code of Conduct

For HDOs that maintain a vendor or third party code of conduct (e.g., for any vendor handling PHI or operating in clinical workflows), organizations should update this document to incorporate AI-specific expectations. This may include:

- Requirements for vendors to maintain responsible AI governance programs.
- Expectations for transparency about how AI is developed, tested, and monitored.
- Requirements for vendors to proactively disclose AI incidents that may affect the HDO.

An updated third party code of conduct sets clear expectations at the start of the relationship and can be referenced in vendor contracts and Request for Proposals (RFPs).

A Note on Cross-Functional Accountability

Third party management is inherently cross-functional and should not be treated as solely a legal or compliance responsibility. Effective third party management and governance requires active engagement from clinical operations, IT, business units, among others throughout the vendor relationship lifecycle. Three functional areas carry distinct and complementary roles: (1) Legal and Compliance typically own contractual language and regulatory alignment; (2) Clinical and IT Operations typically own technical and clinical accountability for how AI is integrated, monitored, and used in practice; and (3) Business Units, including the departments or teams that initiate or use AI solutions, should understand the obligations they carry and actively loop in legal and use case specific governance when issues such as scope creep, unexpected performance, or vendor-driven changes arise. Organizations should define how these functions coordinate, especially for situations that arise outside of formal contract renewal cycles. Smaller organizations without large legal departments should plan for this cross-functional role even with limited resources, potentially by designating a single responsible liaison across functions.

Note: How To Prepare for Vendor Negotiations

Before approaching any vendor about accountability and contract terms, HDOs should achieve internal alignment. Skipping this step is a common reason vendor negotiations stall due to legal, clinical/administrative, IT, etc. arriving at the table with different priorities and vendors exploiting the gap. The following internal alignment steps should be completed first.

1. **Convene the internal negotiation team.** At minimum, this should include a legal or contract representative, a compliance or privacy officer, a clinical or area expert champion who understands how the AI will be used in practice, and an IT representative. Additional staff may be needed based on the use case. Establish a single decision-maker or "owner" for the negotiation.
2. **Agree internally on the must-haves and walk-away criteria.** Before the vendor conversation, the internal team should agree on:
    a. The non-negotiable requirements: clauses you will not execute a contract without
    b. The preferred-but-flexible terms: positions you will advocate for but may accept alternatives on
    c. The walk-away criteria: conditions that would cause you to pause or terminate the engagement.
3. **Assess the HDO's leverage honestly.** The negotiating position may depend on the organization's size, the vendor's market position, the uniqueness of the solution, and whether alternatives exist. Smaller organizations working with large, dominant AI vendors (e.g., major EHR platforms) may have limited leverage on standard contract terms and should plan accordingly by focusing on achievable asks (e.g. disclosure, monitoring obligations, model documentation).
4. **Complete the AI intake to inform Responsible AI (RAI) review.** The information gathered during vendor AI intake (e.g., model card information, intake evidence gathering, training data disclosure, known limitations, regulatory status, etc.) directly informs which contract terms you need. *More information about AI Intake is included below.
    a. There is more than one acceptable way to sequence these steps. HDOs should pick the sequencing that fits their internal governance maturity and the strength of their vendor management discipline. The important takeaway is that intake findings demonstrably shape the final contract, whether they precede signature or drive a subsequent amendment.
5. **Optional: Consider a multi-vendor evaluation (RFP or "bake-off") where feasible.** Where multiple viable vendors exist, organizations have suggested running a structured RFP or side-by-side "bake-off" could meaningfully strengthen the organization's negotiating position. Competing vendors may be more willing to accept disclosure, monitoring, and other components when they know those terms are being used as scored procurement criteria, and comparative performance data gathered during the evaluation can be referenced in later negotiations to justify required SLAs or protections. HDOs should weigh the cost considerations: standing up parallel proofs-of-concept across multiple vendors is resource-intensive, and the return on investment (ROI) of this approach depends on the use case and the number of credible vendors in the market.

Tip: Smaller organizations without dedicated AI legal expertise should reach out to their state hospital association, FQHC network, or regional health information exchange. Many are developing shared AI contract templates and peer networks specifically to help smaller organizations negotiate more effectively.

Keeping the above background context in mind, the below implementation guidance describes how HDOs may consider operationalizing each baseline operational control.

Baseline Operational Control for Subdomain 4.4: Third Party Management

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.

Third Party Management of Disclosure, Roles, and Risk Mitigation

Control 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.

Implementation Guidance

Control 4.4.1 applies to all third party relationships. For the purposes of this control, an 'AI vendor contract' includes any agreement governing the use, access, or integration of an AI solution into the HDO's workflow, including but not limited to vendor contracts, subscription agreements, software licenses, data sharing and use agreements, relevant business associate agreements (BAAs), and data use agreements (DUAs).

Third Party Management Workflow

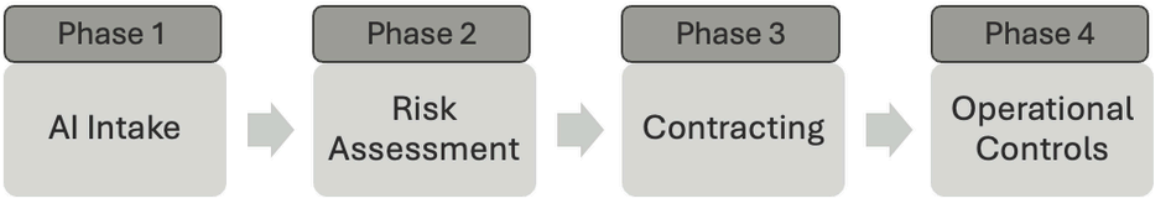

Most organizations move through four phases when bringing an AI solution from interest to in-use governance. They are not strictly linear; later phases regularly feed back into earlier ones, especially when a vendor updates the model, when monitoring surfaces a new risk, or when a use case expands.

Workflow Components

- **AI Intake** gathers vendor attestations and evidence against a common set of AI controls and intake questions. The point of intake is to put a usable evidence base in front of the assessors and to inform risk

assessment and later contracting.

- **Risk Assessment** is where the organization classifies the solution against its own risk tolerance and applies its organizational positions. The risk assessment will impact governance triage, responsible use and lifecycle management, contracting, and operational controls.
- **Contracting** locks in the agreed expectations between the HDO and the vendor. AI addenda are starting points; their substance is shaped by the intake and risk assessment during negotiation. Some HDOs may choose to start with contracting and modify terms following review of the solution by the governance body, while others may choose to start with governance to help guide development of contracting terms.
- **Operational Controls** handle the risks that the contract did not, or could not, fully resolve. These are workflow-level safeguards owned by the deploying organization.

After go-live, monitoring and re-review feed back into intake whenever there is a material change to the model, the use case, or the regulatory picture.

Phase 1: AI Intake

Before beginning intake on any individual AI solution, it may be helpful to start with a working AI inventory of all current AI vendor relationships. This is the foundation for prioritizing which vendors to intake first and identifying existing contractual gaps.

- Conduct a review of existing vendor contracts, software licenses, subscription agreements, BAAs, and data use agreements (DUAs) to identify vendors that provide or enable AI functionality.
- For each AI vendor relationship, document: the name of the AI solution, the ingestion modality (standalone product vs. embedded feature), the type of use case (clinical, operational, administrative), the applicable contract(s), and the contract expiration or renewal date.
- HDOs should pay particular attention to AI-enhanced features within existing EHR, imaging, care-management, or revenue cycle platforms; these are frequently overlooked because the AI is delivered as an update rather than a new procurement.

> **Note**
>
> For smaller organizations, start with the existing EHR vendor. Most major EHR platforms have introduced AI features in recent years. Smaller organizations should review the current agreement to determine what disclosures or governance rights, if any, are already in place.

AI Intake is the first conversation with the vendor and the first organized look at the solution. It is a structured way to ask the same questions of every AI vendor, so that the governance structure responsible for risk assessment, contracting, and operational controls are working from the same evidence. HDOs typically have a 2-part intake process with external vendors (e.g., Part 1: initial, light-touch set of intake questions to triage solutions; Part 2: additional, more robust set of intake questions for higher risk solutions). See the Intake/Procurement section of Subdomain 4.1: Responsible AI Lifecycle Management and Use Playbook for more detail.

Purpose of Phase 1: AI Intake

- Surface what the AI is, what it does, who it is for, and what is known and unknown about its performance, fairness, security, and lifecycle.
- Collect the vendor's attestations and supporting evidence (model cards, validation summaries, certifications, BAA templates, monitoring plans).
- HDOs should produce a single evidence package that the governance structure can reference, without redundant duplicate questionnaires.

Looking Ahead: CHAI AI Intake Questionnaire

Most HDOs have an intake process for new technology solutions. For AI solutions in particular, 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 controls will cover the following domains:

- Clinical Safety & Effectiveness
- Data Integrity & Lineage
- Fairness & Non-discrimination
- Governance Structure & Oversight
- Legal & Regulatory Compliance
- Model Validation & Performance Monitoring
- Security & Access Control
- System Reliability & Resilience
- Transparency & Traceability
- User Experience & Workflow Integration
- Impact Measurement & ROI

The process for arriving at these areas and their accompanying questions/controls include synthesizing feedback from several dozen leading HDOs (more than 800 individual intake questions were collected from approximately 16 questionnaires used by participating health system members).

The primary goal of this standard set of AI intakes 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 would reduce significant burden during the intake and procurement process for both vendors and HDOs. This also allows organizations early in their process to know which questions to ask, and what evidence to look for.

Note

Once an AI intake questionnaire is in place, HDOs should consider embedding it formally into the procurement workflow. This might look like:

- Updating existing vendor intake questionnaires and procurement checklists to include AI-specific questions.
- Requiring AI vendors to provide a model card or equivalent documentation as a condition of contract execution for any AI solution.
- Integrating AI intake findings into contract negotiations: use identified risks or gaps in vendor documentation to inform specific contract requirements.

Looking Ahead: AI Liability in Healthcare

Existing healthcare liability frameworks (medical malpractice, products liability, negligence) were not designed with AI systems in mind, and their application to AI-assisted direct or indirect clinical decisions is actively being litigated and debated. These frameworks often assume there is an identifiable human decision-maker who made a call that could have gone differently; AI distributes that decision across a system in ways none of these frameworks were designed to handle.

This is why prospective contractual risk allocation, not litigation, is currently the most practical tool for managing AI liability in healthcare.

What We Know

- Contractual allocation of risk between HDOs and AI vendors is currently the primary tool available, which makes third party contract review one of the most actionable steps organizations can take now.
- Regulatory bodies (Food and Drug Administration, Office of the National Coordinator for Health Information Technology, Federal Trade Commission) have issued guidance on AI/ML in specific contexts (e.g., Software as a Medical Device), but comprehensive federal AI liability law does not yet exist.

What Remains Unsettled

- Depending on the use case, whether liability for AI-assisted harm rests with the user, the organization, the vendor, or is shared, and how courts will weigh each party's role.
- How "learned intermediary" doctrine applies when AI outputs influence (but do not replace) clinical judgment.
- Whether failure to adopt AI, or failure to override an AI recommendation, could itself constitute negligence in future cases.
- How state-level AI laws (several states have enacted or proposed healthcare AI regulations) interact with federal frameworks.

Recommended Action: Organizations should consult qualified legal counsel when reviewing AI vendor contracts, particularly provisions related to indemnification, limitation of liability, and data ownership.

Key Tools & Resources

CHAI AI Intake Questionnaire

- *A standardized set of AI intake questions organized across 11 domains (clinical safety & effectiveness, data integrity & lineage, fairness & non-discrimination, governance structure & oversight, legal & regulatory compliance, model validation & performance monitoring, security & access control, system reliability & resilience, transparency & traceability, user experience & workflow integration, and impact measurement & ROI) for vetting AI vendors during intake. This document is in draft as of May 2026.*

Phase 2: Risk Assessment

After intake, the organization should assess the solution against its own risk tolerance. The purpose of this phase is to apply the organization's own organizational positions to the intake responses and evidence and produce a risk assessment that drives the next two phases.

AI intake responses and evidence gathered should inform how the HDO conducts their risk assessment. HDOs should review the Subdomain 4.2: Risk and Impact Assessments Playbook for guidance and structure related to risk assessment and related risk stages.

A risk assessment is most useful when it explicitly identifies which risks the organization plans to mitigate, which it plans to transfer, and which it plans to accept. In some cases, it will also determine when an organization chooses not to move forward with a solution, given its risk profile and the organization's risk tolerance. Mitigation through contracting is described in Phase 3. Mitigation through workflow design is described in Phase 4.

Note
The risk assessment phase is an appropriate gate before contracting. Organizations may consider the following:

- Embedding your organization's AI governance review (which may consist of initial AI intake, risk assessment, evaluations) as a mandatory gate in the procurement approval workflow.
- Linking the responsible AI governance review process to contract negotiation, so that any findings, such as model limitations, bias or subgroup performance gaps, or unacceptable risk profiles, can still be used to amend contract terms before execution.
- For organizations using a standard AI intake questionnaire, ensuring the questionnaire is reviewed by legal and/or compliance before contract sign-off.
- If a contract is executed before a full AI governance review can be completed, document this with a defined timeline for completing the review and a process for contract amendment if needed.

Phase 3: Contracting

Risk assessment informs contracting. In some cases, this may be a new contract (especially for new vendors), but in many cases, this may be an AI addendum, appendix, or amendment to an existing master agreement and Business Associate Agreement (BAA). The addendum is the place where the disclosure, monitoring, change notice, data use, and responsibility expectations identified during intake and risk assessment become enforceable obligations between the parties. Note that it is important to check that any addenda do not conflict with terms in the existing contract.

Purpose of Phase 3: Contracting

- Document the expectations the HDO and the vendor have agreed to. This should cover disclosure, data use, change notice, monitoring, and responsibility allocation.
- Create a record of accountability if the solution later fails, drifts, or is materially modified.
- Provide hooks the organization can use later (e.g., audit rights, notification triggers, performance criteria, termination conditions) without dictating every operational detail in advance.

What is an AI Addendum?

An AI addendum or appendix is a self-contained set of AI-specific terms that attaches to a new or existing contract structure. The advantage is that it does not require the organization to renegotiate the underlying master agreement and it sits within the BAA, which already governs PHI handling. The disadvantage is that the addendum's force depends on its order-of-precedence and on how the underlying agreement defines "services," "deliverables," and "confidential information".

Most AI addenda address some combination of:

- (i) What counts as AI for the purpose of the addendum.
- (ii) What notice or approval is required before AI is used.
- (iii) Data use restrictions.
- (iv) Ownership of outputs and derivative data rights.
- (v) Representations and warranties.
- (vi) Monitoring and incident reporting.
- (vii) Change management and material modifications.
- (viii) Subcontractor and foundation-model disclosure.
- (ix) Audit and reporting rights.
- (x) Specification of lifecycle management, operational, and regulatory responsibilities across the parties.
- (xi) Disclosure of known model limitations and risks.
- (xii) Termination and sunsetting.

Two example addenda, one from Mercy and one from Kaiser Permanente, are reproduced in the Tools & Resources section of this playbook. They are different in posture and substance, and together they illustrate the range of approaches in the field.

Specification of lifecycle management, operational, and regulatory responsibilities across the parties is a key component of the contract, as defined in **control 4.4.1**. Some areas that are important to discuss between organizations and vendors include: pre-deployment validation, staff training, monitoring, incident reporting/feedback channels, regulatory submissions, and termination or sunsetting processes. Where responsibilities are shared, the contract should describe how coordination will occur. This specification should account for the full AI solution or system – including the model, the technical infrastructure, and the human workflows in which the AI operates – and should align with any applicable regulatory requirements (e.g., FDA guidance on SaMD, ONC regulations related to decision support interventions). Below are some example questions that HDOs may want to discuss with vendors.

- **Lifecycle Management Responsibilities** (examples): Who is responsible for validation of the AI solution prior to deployment or monitoring it post-deployment? Who oversees clinical, operational, or administrative staff training on appropriate use? Who defines and monitors performance and use? Who manages decommissioning or modifications if performance degrades? Who ensures that end users understand the scope of their clinical or administrative accountability when relying on AI outputs?
- **Operational Responsibilities** (examples): Who manages system integration and technical infrastructure? Who is the internal owner of the AI vendor relationship on a day-to-day basis? Who tracks model version updates and coordinates with IT for change management? Who handles scope changes requested by the vendor during or after the pilot phase?
- **Regulatory Responsibilities** (examples): When relevant, Who is responsible for determining whether the AI solution qualifies as FDA-regulated Software as a Medical Device (SaMD)? Who owns FDA Mandatory Device Reporting (MDR) submissions in the event of an adverse AI-related event (where applicable)? Who monitors ONC clinical decision support regulations applicable to the solution? Who owns patient notification obligations if a solution is found to have caused harm?
- **Incident Reporting Responsibilities** (examples): When an AI solution or system contributes to an adverse event, the contract should specify:
  - Which party owns internal event documentation.
  - Who is responsible for FDA MDR reporting if the solution is classified as SaMD.
  - Which party should initiate patient notification.
  - The required timeline for each obligation. This is a consequential gap in many current vendor agreements and should be addressed explicitly.
- **End User Transparency and Education** (examples): When a third party AI solution is used in direct workflows or patient interactions, the HDO should ensure that end users (clinical, operational, or

administrative staff, and where appropriate, patients) are made aware that a third party AI solution is involved.
    - Transparency disclosures may include: the nature of the AI solution and its role in the workflow; the fact that a third party vendor is involved; known limitations that are relevant to users; and how to report unexpected behavior.
    - Education and training of end users is an important risk-mitigator. See Subdomain 4.5: Education, Training, and Feedback Playbook for further guidance on operationalizing end user education.

Key Tools & Resources

Example Template: Assign Roles and Responsibilities

- *This document includes a sample set of responsibilities that healthcare delivery organizations (HDOs) may build on during vendor contract discussions to clarify role ownership across lifecycle management, operational, and regulatory responsibilities.*

Disclosure of Known Model Limitations and Risks is another key component of the contract, as defined in **control 4.4.1.** Vendors should be contractually obligated to disclose known limitations, risks, and performance characteristics of their AI solutions or systems — including subgroup performance variability, out-of-distribution behavior, conditions under which the model should not be used, and any known safety concerns identified during validation. This disclosure may take the form of a model card, technical documentation, and/or a contractually specified disclosure schedule. Disclosure obligations should be ongoing: vendors should be required to notify the organization when material changes to the model occur, including version updates, retraining events, or newly identified limitations. Vendors should also be contractually obligated to notify HDOs if the vendor's FDA classification changes (e.g., moves into SaMD). Regulatory reclassification needs to be an explicit disclosure trigger.

Note: Foundation Models and Open-Source Components

For vendor-built AI solutions that incorporate third party foundation model (e.g., large language models licensed from external providers) or open-source components, the vendor should provide at intake:

(a) The identity of the third party model provider.

(b) The license terms under which the component is used and any restrictions relevant to healthcare data.

(c) Any vendor-applied constraints on the foundation model's behavior specific to the healthcare deployment context.

(d) The vendor's process for monitoring and disclosing changes made by the foundation model provider that affect system behavior.

HDOs should treat absence of this documentation as equivalent to undisclosed model limitations. HDOs should factor the depth of available documentation into their risk processes (more information about this in Subdomain 4.2 Risk and Impact Assessments Playbook).

Looking Ahead: Additional Disclosure Requirements for Agentic and Autonomous AI

For solutions that operate autonomously or as agents (i.e., the AI can take actions across other tools or systems on the user's behalf with limited or asynchronous oversight), the vendor should additionally disclose:

- **Scope of autonomy**: Which actions can the agent take without confirmation? Which require human approval?
- **Tool and integration access**: Which systems can the agent read from, write to, or invoke? Which external APIs or third party services does it depend on?
- **Action reversibility**: Which actions are reversible and which are not (e.g., medication orders, billing submissions, external communications)?
- **Escalation and circuit breakers**: What conditions trigger human escalation or automatic halt?
- **Memory and context boundaries**: What persistent memory does the agent retain across sessions, and what mechanisms exist to scope or expire it?
- **Multi-agent considerations**: If the system orchestrates other agents, each should be enumerated and governed.
- **Audit trail granularity**: Logs sufficient for incident reconstruction.

Note about Post-Deployment Monitoring as a Contractual Requirement

Post-deployment monitoring is a distinct and critical contractual requirement. Vendor contracts should include an explicit monitoring plan specifying:

- What performance metrics the vendor will track and report.
- The frequency and format of monitoring reports shared with the HDO.
- The vendor's obligations when performance degradation, model drift, or safety signals are detected.
- The HDO's right to request monitoring data. HDOs should be aware that vendors may resist ongoing monitoring obligations.
- The consequences for missed reporting deadlines or inadequate data. Without defined consequences, monitoring obligations are frequently treated as aspirational/optional. Contracts should include concrete remedies such as service credits, fee withholding or refund rights, expanded audit rights, and, for repeat or sustained failures, suspension of new feature activation or new use cases pending remediation.
- A direct link between monitoring compliance and contract renewal. Renewal should be conditioned on the vendor's demonstrated record of meeting reporting obligations during the prior term, and KPI performance over the term should be formally considered in the renewal decision.

Service Level Agreement (SLA) Considerations

SLA discussions in AI contracts often default to generic uptime targets and miss the AI-specific considerations that determine real world reliability. SLAs should also address:

- **Down-time planning and recovery objectives:** Recovery Time Objective (RTO) and Recovery Point Objective (RPO) defined per AI service, with a graduated approach for critical patient safety solutions.
- **Business continuity:** Defined manual workarounds for each clinical or operational workflow that becomes degraded if the AI is unavailable; the contract should recommend the vendor surface and document these workarounds.
- **Disaster recovery:** These may include geographic redundancy, backup model versions, failover procedures, and the vendor's tested recovery cadence. Vendor should provide recurring disaster recovery summaries.
- **AI dependency map:** The vendor should provide a written map of upstream dependencies (e.g., foundation models, hosted inference services, cloud regions, identity providers, data lakes, and third party APIs) and identify what functionality becomes unavailable if each dependency goes offline.

- **Notification timelines for outages and dependency failures:** Including outages of upstream providers (e.g., the cloud or foundation model vendor) not owned directly by the AI vendor.

These items should sit alongside KPIs and performance metrics in the SLA.

Note: Standard Vendor Management Practices (KPIs and Accountability)

AI vendor relationships benefit from the same foundational vendor management discipline applied to any critical vendor. Contracts and relationship governance should include:

- **Defined KPIs up-front** in the contract or statement of work (e.g., model performance thresholds, incident response times, update cadence, training/support responsiveness).
- **KPIs tied to Service Level Agreements (SLAs)** with defined remedies such as service credits, fee reductions, and/or termination rights.
- **A recurring accountability cadence** at which KPI performance is reviewed, issues are escalated, and corrective actions are documented. This should occur separately from contract renewal
- **Documented KPI performance as input to contract renewal**, so that renewal decisions are grounded in evidence. Separating KPI review cadence from contract renewal ensures that multi-year contracts don't weaken KPI enforcement.

These standard vendor management practices help ensure that AI-specific contractual protections are operationalized over the life of the engagement.

Note: FDA SaMD Considerations

When a third party AI solution may qualify as FDA-regulated Software as a Medical Device (SaMD), responsibility allocation in the vendor contract becomes especially consequential. HDOs should:

- Ensure the contract clearly identifies whether the solution has been or may require FDA clearance or approval.
- Specify which party is responsible for pursuing, maintaining, or reporting changes to FDA authorization.
- Address how responsibility shifts if the tool is modified in ways that change its regulatory classification.
- Include specific obligations for adverse event reporting to the FDA under 21 CFR Part 803. Legal counsel familiar with FDA SaMD regulations should review any contracts involving tools that may fall under this framework.

When a Contracted Vendor Fails to Meet Obligations: Post-deployment Enforcement

It is important to establish a defined path for what happens when a vendor fails to meet contractual obligations after signing. Contracts should include an explicit tiered enforcement mechanism with defined timelines:

1. **Tier 1: Notice and Cure**: formal written notice of the deficiency, the specific contractual obligation implicated, and a defined cure period during which the vendor should remedy the issue.
2. **Tier 2: Remediation Plan**: if the deficiency is not cured, the vendor should submit and execute a written remediation plan with milestones, owners, and acceptance criteria; the HDO has the right to reject an inadequate plan.
3. **Tier 3: Access Suspension**: if remediation milestones are missed, the HDO may suspend activation of new features, new use cases, or additional data access; and/or withhold fees until the vendor returns to compliance.
4. **Tier 4: Termination for Cause**: with defined transition assistance obligations, data return or destruction requirements, and continued cooperation on open incidents.

Each tier should have pre-defined triggers, decision-makers, and timelines documented in the contract so enforcement does not depend on ad hoc escalation.

Key Tools & Resources

Example Template: Service Level Agreement (SLA) Enforcement Tiering

- *Procurement teams, legal counsel, and AI governance leads can use this template as a starting anchor when negotiating AI vendor SLA enforcement. Adapted from model-clause patterns referenced by Health Sector Coordinating Council (HSCC).*

Vendor Failure: Bankruptcy, Acquisition, Shutdown

Tiered enforcement assumes a vendor that exists and is operational. Contracts should also plan for cases in which the vendor itself fails. HDOs should consider including the below concepts in AI vendor contracts:

- **Bankruptcy or insolvency triggers:** Filing of bankruptcy, assignment for the benefit of creditors, or similar action(s) should trigger access to (a) source code escrow or operating documentation, (b) accelerated data return / destruction obligations, and (c) the right to terminate without further notice.
- **Change-of-control or acquisition provisions:** Written notice to the customer of any proposed change of control, with the customer's right to (a) re-review the deal for AI governance compatibility (e.g., new owner's data practices, new owner's foundation model dependencies), (b) terminate without penalty if AI governance posture materially worsens, and (c) require the acquirer to ratify the AI addendum in writing.
- **Product sunset or discontinuation:** Defined notice period and required transition support, including export of customer data and any organization-specific configurations in usable form.
- **Source code and model escrow:** For high-risk or critical patient safety solutions, require escrow of model weights, inference code, documentation, and/or runbooks, with release triggers tied to vendor failure events.
- **Support disappearance:** If the vendor ceases supporting the product (even without insolvency), the contract should specify that the customer may continue using the most recent version (e.g., under a perpetual, royalty-free license for a defined wind-down period) and that the vendor cannot remotely disable the product without honoring the wind-down terms.

Note: Integrating AI into Business Associate Agreements (BAAs)

For AI solutions that involve protected health information (PHI), organizations should ensure that a BAA with the vendor includes an AI-specific appendix or addendum that covers components of this playbook. This AI-specific appendix or addendum should be treated as a living document and updated when the AI solution is materially modified, when the scope of PHI access changes, or when new risks are identified. HDOs should not assume that standard BAA language adequately addresses AI model disclosure, roles and responsibilities pertinent to the AI system and known risk, or contractual and operational risk mitigation provisions. A dedicated AI appendix or addendum provides an important mechanism to document AI-specific obligations within an existing, familiar contractual structure.

Phase 4: Operational Controls

Operational controls are the workflow-level safeguards the HDO puts in place to mitigate risks that the contract did not, or could not, fully resolve. They are owned by the deploying organization, not the vendor, and they are how residual risk is managed in practice. They become especially important when the vendor will not accept a

particular contractual obligation, when the risk lives in human workflow rather than in the model, or when the AI is embedded in a platform the organization has limited contractual leverage over.

Purpose

- Manage the risks that survive contracting, which include risks at the workflow, training, and oversight level.
- Provide a stable mechanism for human oversight, escalation, and rollback that does not depend on the vendor.

Common Operational Control Patterns

Included below are examples of operational controls that HDOs may consider. These examples are non-exhaustive and shared for illustrative purposes.

- Consent and disclosure at activation: Patient (or end user) consent obtained before the solution is used; consent captured into the record in plain language rather than as boilerplate.
- Default-off for sensitive contexts: The solution is disabled by default for encounter types or workflows where the risks outweigh the benefits, with a documented procedure to enable it on a case-by-case basis.
- Human review of outputs: for certain use cases, AI outputs are presented as drafts; sign-off requires affirmative human action; override and edit rates are tracked.
- Named owner and escalation path: A clinical informatics, IT, or operational owner is identified for certain AI solutions, with documented contact information and vendor-side counterparts.
- Independent rollback procedure: The HDO can disable the solution independently of the vendor for a safety or privacy concern.
- Configuration to minimum data exposure: Tenant settings are configured to the lowest retention, lowest data sharing, and tightest residency options the vendor permits.
- Periodic re-review: At a defined cadence and on triggers (material modification, foundation model change, incident).
- AI-tagged incident reporting: Incidents are captured by the HDO with an AI tag for review.
- Agent-specific guardrails (for autonomous solutions): Develop action allow-lists, environment isolation, dry-run modes for high-impact actions, and per-action audit logs.

Implementation Guiding Questions

Below are standard questions that HDOs may ask to guide discussion for **control 4.4.1**:

- Have you reviewed your existing vendor contracts and BAAs to identify whether AI-specific clauses, such as model disclosure and responsibility, are already present (or absent)?
- Do you have a process for distinguishing between AI solutions that require new contracts and those introduced as AI-enhanced features within existing vendor systems (e.g., EHR-embedded AI)?
- What internal stakeholders (legal, compliance, clinical informatics, IT, clinical operations) need to be involved in reviewing and approving AI-specific contract language?
- How will you handle situations where a vendor refuses to accept responsibility for certain risks or declines to disclose limitations? Do you have escalation criteria for walking away from a contract?
- Is your organization prepared to require disclosure of model cards or equivalent documentation as a condition of contract execution, and how will you evaluate what vendors provide?
- Have you defined which roles within your organization are responsible for ongoing monitoring of AI vendor compliance with contract terms after the solution is deployed?
- Do your contracts specify conditions that would trigger re-negotiation, performance review, or contract termination (e.g., model updates, safety events, regulatory changes)?
- Do you have a mechanism in place (e.g., contractually with vendor or otherwise) to be immediately notified if an AI solution's FDA classification changes (e.g., moves into SaMD)? Regulatory reclassification needs to be

an explicit disclosure trigger.

- Are relevant clinical, operational, or administrative staff aware of the AI capabilities embedded in solutions they use? Do they understand their accountability (if any) under current contract terms?
- As you build or update AI terms / contract language, how are you ensuring consistency across departments? Do you have a central repository for approved clause templates?
- Have you run an inventory of AI features embedded in your existing enterprise platforms, and re-read each platform's data processing terms for AI model training provisions?
- Have you addressed vendor failure concerns (bankruptcy, acquisition, shutdown) in your standard contract language?
- For agentic / autonomous solutions, have you documented action scope, reversibility, escalation, and audit-trail expectations in the contract?

Illustrative Examples

The following illustrative examples describe how HDOs of varying size and complexity may approach third party management. These examples are illustrative, grounded in literature and common patterns observed across the field, and are not drawn from any specific documented organization.

Illustrative Example #1: Community Health Center

Scenario: A Community Health Center (CHC) in a rural underserved area begins using an AI-powered chronic disease management tool embedded within their EHR platform. The tool surfaces risk alerts for patients with uncontrolled diabetes. When a care coordinator asks the vendor for documentation on how the model was developed and what populations it was validated on, she is told the information is proprietary. The CHC's part-time compliance manager is tasked with reviewing the existing EHR contract and finds no AI-specific provisions. Working with a regional primary care association, the CHC adopts an AI Appendix template developed by a national CHC network. The appendix requires the EHR vendor to provide a model card within 60 days, disclose any model updates that affect clinical outputs, and acknowledge that the CHC retains oversight authority over how the tool is used in its workflows. Upon renewal of the EHR agreement, the appendix is formally incorporated.

Alignment to Subdomain 4.4: This example demonstrates how a resource-constrained organization can leverage shared resources (regional association templates, peer network guidance) to meet third party management requirements. The focus on amending an existing BAA rather than creating a new agreement reflects the playbook's principle of integrating AI governance into existing infrastructure.

Illustrative Example #2: Medium-Size or Regional HDO

Scenario: A 250-bed community hospital system in the Midwest deploys an AI-assisted sepsis early warning system from a third party vendor. Prior to contract execution, the HDO's clinical informatics team conducts a two-part intake review. Part 1 identifies the solution as High Risk (clinical decision support, patient safety-critical). Part 2, completed jointly with the vendor, reveals that the model was validated primarily on a patient population different from the hospital's demographic mix. The legal and compliance team negotiates contract language requiring: (1) the vendor to provide updated model cards within 30 days of any retraining or model update; (2) a joint performance review at 6 and 12 months post-deployment; and (3) a shared responsibility matrix specifying that the vendor is accountable for model accuracy and the hospital is accountable for clinical workflow integration and clinical, operational, or administrative staff training. The shared responsibility matrix is incorporated into the hospital's clinical operations (CONOPS) for the sepsis program.

Alignment to Subdomain 4.4: This example demonstrates how a medium-sized hospital can apply risk-based contracting using risk categorization from the intake process to determine the depth of contractual protections required. The joint performance review and shared responsibility matrix reflect the recommendation to specify lifecycle management, operational, and regulatory responsibilities.

Illustrative Example #3: Large HDO

Scenario: A large regional health system with 12 hospitals and over 3,000 employed clinicians establishes a formal AI Governance Committee that includes representation from legal, compliance, IT, clinical operations, and patient safety. The committee develops a standardized AI Vendor Contract Addendum and deploys it as a required attachment to all new AI-related vendor agreements system-wide. The AI addendum includes: mandatory model documentation (model card or equivalent, updated annually); a change management notification requirement (30-day advance notice for material changes); and a data access provision allowing the health system to audit AI performance metrics on a defined schedule. When a major EHR vendor announces an AI-powered clinical documentation feature as part of a routine software update, the AI Governance Committee invokes the change notification clause and requires the vendor to submit updated documentation before the feature is activated. The committee's review identifies a documentation gap (e.g., the vendor has not disclosed training data sources) and the feature's activation is delayed pending remediation.

Alignment to Subdomain 4.4: Some large HDOs have the leverage and internal capacity to establish proactive, system-wide AI governance infrastructure. This example demonstrates how a standardized contract AI addendum can serve as a durable governance mechanism across a complex enterprise. The EHR feature scenario reflects a documented and growing concern in the field: AI capabilities being introduced through routine software updates without explicit review. The HDO's ability to delay activation pending documentation review is only possible because the contractual right to require that documentation was established in advance.

Tools & Resources

Included below are tools and resources that may be helpful for developing and operationalizing your own third party management processes.

CHAI AI Intake Questionnaire

- *A standardized set of AI intake questions organized across 11 domains (clinical safety & effectiveness, data integrity & lineage, fairness & non-discrimination, governance structure & oversight, legal & regulatory compliance, model validation & performance monitoring, security & access control, system reliability & resilience, transparency & traceability, user experience & workflow integration, and impact measurement & ROI) for vetting AI vendors during intake. This document is in draft as of May 2026.*

CHAI Partner Ecosystem

- *Collaborate with trusted partners delivering excellence in AI governance, testing, data management, and advisory services - all aligned with the CHAI AI Governance Framework.*

CHAI Applied Model Card (AMC)

- *A standardized documentation framework that describes an AI solution in the context of a specific health use case. Used during procurement to confirm adherence to responsible AI principles.*

Concept of Operations (CONOPS) template

- *Include AI responsibility allocation in Concept of Operations (CONOPS) documents to define accountability across development, deployment, and monitoring. This clarifies which party (vendor vs. healthcare delivery organization) owns each lifecycle obligation.*

Example Template: Assign Roles and Responsibilities

- *This document includes a sample set of responsibilities that healthcare delivery organizations (HDOs) may build on during vendor contract discussions to clarify role ownership across lifecycle management, operational, and regulatory responsibilities.*

AI Terms and Standardized Contract Language

- *AI Terms from Mercy and Kaiser Permanente are included below. CHAI recommends organizations develop or adopt a reusable clause library for AI-specific BAA appendices or addenda.*

AI Terms from Mercy *(verbatim)*

The text below are standardized AI terms developed by Mercy that organizations may reference as an appendix or addendum to existing vendor contracts or BAAs. Inclusion is for illustrative reference only; consult counsel before adapting any language.

ARTIFICIAL INTELLIGENCE TERMS – SHORT FORM

Service Provider shall not use artificial intelligence, machine learning, generative AI, or similar technologies (collectively, "**AI**") in the performance of the Services without first providing prior written notice to Customer and obtaining Customer's prior written approval to the extent such AI: (a) generates, recommends, or automates any action or decision to be taken by Customer or Customer personnel; (b) generates outputs that are incorporated into, relied upon for, or submitted in connection with billing, coding, reimbursement, or payor communications involving Customer; (c) interacts directly with patients or consumers, or generates content or communications delivered to patients or consumers; (d) supports or informs clinical care, scheduling, triage, prioritization, or access to care; (e) processes Personal Data other than in a purely passive or technical capacity; or (f) supports or informs regulatory, compliance, audit, or reporting obligations attributable to Customer. Notwithstanding the foregoing, Service Provider may use AI without prior approval solely for its internal administrative or back-office operations provided that all of the following conditions are met: (i) the AI does not access, process, analyze, or use any Customer Data, including Patient Information or data derived from the Services; (ii) the outputs, recommendations, or results of the AI are not provided to, shared with, relied upon by, or incorporated into any Services, Deliverables, communications, or other materials provided to or accessible by Customer; (iii) the AI does not fall within any of clauses (a)-(f) above; and (iv) the AI is used exclusively for functions such as Service Provider's internal accounting, human resources, facilities management, or similar back-office purposes. Customer may condition any approval under this Section on the parties' execution of additional terms governing the use of AI.

For reference:

"**Customer Data**" means all data, records, content, or other information (including Personal Data) that is provided, made available, or otherwise originated by or on behalf of Customer, its Affiliates, patients, or customers, including but not limited to patient, clinical, operational, financial, or administrative data, regardless of format or medium. Customer Data includes any aggregations, de-identified data, models, reports, results, or other outputs generated from or based upon such data, whether generated by the Service Provider, Customer, or any third party. For clarity, Customer Data shall be treated as Confidential Information under this Agreement, provided, however, that to the extent Customer Data constitutes PHI, such PHI shall be subject to the BAA.

"**Personal Data**" means any information that may be processed at any time by Service Provider in connection with or incidental to the performance of the Agreement that: (a) relates to an identified or identifiable individual; (b) identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular individual, device, or household; or (c) any other data identified as personally identifiable information or similar protected data under data protection laws. Information Processed by Service Provider deriving from the foregoing is also considered Personal Data.

ARTIFICIAL INTELLIGENCE TERMS – LONG FORM

1. **Definitions**. Capitalized terms not otherwise defined under the Agreement shall have the following definitions:

(a) "**AI Technologies**" means deep learning, machine learning, and other artificial intelligence (AI) technologies, including, but not limited to, algorithms, software, and systems that make use of or employ neural networks, natural language processing, computer vision, statistical learning algorithms, reinforcement learning, or generative models.

(b) "**Covered AI Technologies**" means AI Technologies that influence, recommend, automate, or otherwise affect decisions, actions, outputs, or outcomes involving any of the following: (a) clinical care, including diagnosis, treatment, or medical ordering; (b) triage, scheduling, prioritization, or access to care; (c) patient or consumer interactions (e.g., chatbots, symptom checkers); (d) billing, coding, documentation, utilization management, or payor communications; (e) employment activities, including hiring, screening, performance evaluation, or compensation decisions; (f) data aggregation, summarization, or synthetic data generation; or (g) regulatory or compliance obligations under federal or state law.

(c) "**Material AI Modification**" means any update, retraining, reconfiguration, integration, or other change that materially alters the functionality, intended use, outputs, data inputs, decision-making role, or system architecture of AI Technologies. This includes, but is not limited to: (a) changes to model architecture, algorithms, or learning methods; (b) modifications to data pipelines or training, validation, or inference datasets; (c) retraining or fine-tuning of models; (d) updates to thresholds, scoring logic, or decision criteria; (e) integration with new external systems, APIs, or data sources; (f) introduction of new features or capabilities that expand scope, functionality, or risk; or (g) any change that may materially affect the safety, accuracy, bias, fairness, explainability, auditability, or regulatory classification. For the avoidance of doubt, a change shall constitute a Material Modification regardless of whether it is accompanied by a version number or product name update.

2. **Pre-Deployment Requirements**. Prior to the implementation, enablement, or Material AI Modification of Covered AI Technologies in the provision of Services, Service Provider shall: (a) provide written notice to Customer describing the nature, scope, purpose, and functionality of the Covered AI Technologies; (b) supply a comprehensive model card or equivalent documentation sufficient to evaluate the safety, efficacy, fairness, explainability, transparency, data sources, training methodologies, validation results, known limitations, potential biases, and ongoing governance of the Covered AI Technologies, in a form and with content reasonably satisfactory to Customer; and (c) obtain Customer's prior express written consent, which may be provided through a mutually executed written addendum or amendment hereto setting forth additional terms, conditions, limitations, and oversight mechanisms related to the Covered AI Technologies. Nothing in this Section shall be construed to require Service Provider to disclose trade secrets or other proprietary information that is not publicly available.

3. **Ongoing** Compliance. At any time Covered AI Technologies are active, deployed, or in use, Service Provider shall: (i) maintain and operate a system to continuously test, validate, and monitor the performance, accuracy, safety, and behavior of the Covered AI Technologies, including monitoring for performance degradation, model drift, data quality issues, and risks such as hallucinated outputs, incorrect recommendations or decisions, discriminatory outcomes, unfair bias, harmful results, privacy violations, or non-compliance with Applicable Law, and provide Customer with supporting evidence (including logs, audit trails, test results, performance metrics, and incident reports) upon request; and (c) notify Customer in writing within two (2) business days of identifying or becoming aware of any performance degradation, adverse incident, regulatory inquiry, or other risks described in subsection (i), along with a detailed assessment of the issue and proposed remediation plan;

and (iii) maintain comprehensive documentation of all Covered AI Technologies, including data provenance, model versions, training procedures, validation methodologies, deployment configurations, and change history, and make such documentation available to Customer upon request. Customer may, in its sole discretion and at any time, suspend, restrict, modify, or terminate use of any Covered AI Technologies pending governance review, upon identification of risks or concerns, for non-compliance with this Section, or for any other reason Customer deems appropriate to protect its interests, and may require remediation, additional safeguards, enhanced monitoring, third-party audits, or other conditions as a prerequisite to continued use. Service Provider shall not charge any additional fees for compliance with this Section or for suspension, modification, or remediation of Covered AI Technologies. Nothing in this Section shall be construed to require Service Provider to disclose trade secrets or other proprietary information that is not publicly available.

4. **Representations and Warranties**. Service Provider represents and warrants that: (a) all AI Technologies comply with Applicable Law; (b) all data used to train, validate, or operate the AI Technologies was lawfully obtained and used, and Service Provider has all necessary rights to such data; (c) the AI Technologies have been tested for bias and discriminatory outcomes, and Service Provider has implemented reasonable controls to identify and mitigate such risks; (d) the AI Technologies perform in accordance with the accuracy, precision, and performance specifications set forth in the applicable SOW or documentation; (e) Service Provider maintains appropriate governance, testing, and quality assurance processes for the AI Technologies, including monitoring for performance degradation and model drift; and (f) the AI Technologies include appropriate

AI Terms from Kaiser Permanente *(verbatim)*

The text below is reproduced verbatim from Kaiser Permanente's Artificial Intelligence (AI) Requirements for Vendors, Contractors and Suppliers (Version 09/3/2025, document reference #1807156v1). Inclusion is for illustrative reference only; consult counsel before adapting any language.

ARTIFICIAL INTELLIGENCE (AI) REQUIREMENTS FOR VENDORS, CONTRACTORS AND SUPPLIERS

Kaiser Permanente uses artificial intelligence ("AI") to improve high-quality and affordable care with consideration for its members, clinicians, and employees and aims for the AI tools that it uses to be safe, reliable, equitable, and trustworthy. To meet this objective, Kaiser Permanente has created these Artificial Intelligence Requirements ("AI Attachment") to set forth the AI requirements for vendors, contractors and suppliers that will support its responsible AI approach.

1. DEFINITIONS

The following definitions apply, unless equivalent definitions are in the Agreement:

1.1. "AI Services" means the artificial intelligence software, products, technology, systems, and/or services, together with all updates and workarounds, corrections, modifications, and improvements provided by Supplier, and as (i) described in an applicable Agreement or otherwise provided by Supplier to Customer and KP Affiliates, or (ii) services, functions, processes, and responsibilities that are required for or inherent in the proper performance and delivery of the AI Services.

1.2. "Agreement" means any agreement between Supplier and a Customer, including any SOW, Purchase Order, or ordering document, under which Supplier is selling, licensing, leasing or supplying Cloud Services, Software, Products and/or Services to a Customer.

1.3. "Confidential Information" for Customer and KP Affiliates, includes, without limitation, Outputs, KP Data, Client Data, protocols, guidelines, workflows, models, policies, processes, algorithms, and input and prompts of Customer and KP Affiliates and their respective contractors, suppliers (other than Supplier), agents, and vendors, consultants, providers, and modifications, enhancements, and derivative works thereto and therefrom regardless of the party that created them. All references to Customer's Confidential Information, includes that of KP Affiliates.

1.4. "Customer" means any KP Affiliate purchasing or licensing products and/or services from Supplier under the Agreement.

1.5. "Cloud Services" means Supplier hosted software, platform and/or infrastructure available for Customer's use through remote access service, including software applications, mobile applications, platform, infrastructure, content, hosting, management, support and maintenance services, together with all relevant documentation, updates and workarounds, corrections, modifications, and improvements provided by Supplier pursuant to the Agreement.

1.6. "Intellectual Property Rights" definition includes all intellectual property or other proprietary rights worldwide, including patents, copyrights, trademarks, service marks, trade names, domain name rights, know-how, moral rights, trade secrets and all other intellectual or industrial property, including all associated applications, registrations, renewals and extensions of those rights.

1.7. "KP Affiliate" means an entity participating in the integrated health care delivery system doing business as Kaiser Permanente®, including Kaiser Foundation Health Plan, Inc., Kaiser Foundation Hospitals, the Permanente Medical Groups,

and all subsidiaries and successors of the foregoing, and Risant Health, Inc., and its subsidiaries and successors. The foregoing entities are collectively, as a whole, referred to as "KP Affiliates".

1.8. "KP Data" or "Client Data" or "Customer Data" definition includes inputs made by Customer's authorized users or Customer or any KP Affiliate into or submitted to the AI Services.

1.9. "Output" means outputs, information, results and other similar materials prepared or generated for or by Customer or KP Affiliates as part of the AI Services, and any modifications, enhancements, or derivative works thereto.

1.10. "Personal Information" or "PII" means personally identifiable information, data or records relating to or concerning any patient, member, plan participant, employee, or contractor of Customer and KP Affiliates, including PHI, employee records and, if applicable, Cardholder Data as defined under the Payment Card Industry data security standards.

1.11. "PHI" means member records and other Protected Health Information under HIPAA.

1.12. "Pre-Existing Material" means any pre-existing data, content or materials owned or licensed by Supplier prior to the start date of the Agreement or application of this AI Attachment, whichever is the earlier, or any data, content or materials independently developed by Supplier outside of the scope of the Agreement, without reference to Customer's Confidential Information.

1.13. "Product" means an item identified under the Agreement that Supplier is offering for sale to KP Affiliates, including equipment, hardware, spare parts, updates, upgrades, related documentation, and included software programs and provided by Supplier pursuant to the Agreement.

1.14. "Risant Health" means Risant Health, Inc., and its subsidiaries and successors.

1.15. "Services" means the services described in the Agreement and provided by Supplier pursuant to the Agreement.

1.16. "Software" means on-premises software program(s) listed and described in the Agreement, including any bug fixes, patches, updates or new releases of the Software, provided by Supplier pursuant to the Agreement.

1.17. "Specifications" definition includes AI Services documentation and any applicable model card provided by the Supplier.

1.18. "Supplier" means a vendor, contractor or supplier providing Cloud Services, Products, Software and/or Services to a Customer.

2. AI SERVICES AND ADDITIONAL AI TERMS

2.1. **Provision of AI Services**. The requirements of this AI Attachment apply to any Supplier providing Cloud Services, Products, Software and/or Services to Customer and KP Affiliates which use and/or incorporate AI Services. This AI Attachment is hereby incorporated by this reference into any Agreement and this AI Attachment takes precedence and governs, unless the parties have agreed in writing to alternative terms and conditions regarding AI Services that expressly take precedence over this AI Attachment. Each reference to "Services," "Cloud Services," "Products," "Software," and related services in the Master Agreement will be deemed to also include the "AI Services," and all of the terms applying to Cloud Services, Services, Products, Software, and related services in the Agreement are deemed to also apply in full to AI Services. References to a SOW, Purchase Order, or ordering document in the Agreement is a reference to the applicable SOW, Purchase Order, or ordering document to which the AI Services relate.

2.2. **Data Usage and Restrictions**. Supplier is expressly prohibited from using Customer's Confidential Information, KP Data, Outputs, or Personal Information (including any anonymized, de-identified or aggregated Customer's Confidential Information, KP Data, Outputs, or Personal Information) to train, fine tune, improve, or perform research and development related to the Cloud Services, Services, Software, and Products, including AI Services, or any other products and services. Supplier is prohibited from de-identifying, re-identifying selling, distributing, commercially exploiting, aggregating, data mining, analyzing, benchmarking, or otherwise using or disclosing any Customer's Confidential Information, KP Data, Outputs, or Personal Information (including any anonymized, de-identified or aggregated Customer's Confidential Information, KP Data, Outputs, or Personal Information) for any purpose other than to provide the Cloud Services, Software, Products and Services, including AI Services, to KP Affiliates under the Agreement.

2.3. **Intellectual Property**. As between Customer and Supplier, Customer will be the owner of all Outputs and all Intellectual Property Rights therein, exclusive of any Pre-Existing Material of Supplier or third party materials contained therein. To the extent that any rights in any Outputs do not vest in Customer pursuant to this provision, Supplier hereby irrevocably assigns in perpetuity to Customer: (i) all worldwide right, title, and interest in and to such Outputs, exclusive of any Supplier Pre-Existing Material or third party materials contained therein, and (ii) all claims, demands and rights of action, both statutory and based upon common law, together with the right to prosecute such claims, demands and rights of action in Customer's own name. Except with the prior written consent of Customer, Supplier will not disclose, use, license, sell or otherwise transfer all or any part of any Outputs to any third party. If Supplier includes any Supplier Pre-Existing Material in any Outputs, Supplier grants to Customer a nonexclusive, worldwide, irrevocable, fully paid, royalty free, perpetual license and right to use, modify, access, perform, execute, display, reproduce, distribute, enhance, and create derivative works of such Supplier Pre-Existing Material in conjunction with Customer's and KP Affiliates' use of such Outputs.

2.4. **AI Representations and Warranties**. In addition to the representations and warranties set forth in the Agreement, Supplier represents, warrants, and covenants that: (a) AI Services, AI Training Data, and Outputs as provided herein, will comply with and will not violate or contravene any applicable law promulgated by any government or regulatory body; (b) Supplier has disclosed to Customer, all artificial intelligence technology it has embedded in or provided with the Cloud Services, Software, Products and/or Services; (c) Supplier has and shall maintain any necessary rights, consents, authorizations, permissions, certifications, and licenses to: i) develop and provision the AI Services, ii) use any data, information, images, videos, code, or materials ("AI Training Data") to train, validate, develop, or test any AI model, tool, algorithm, service, application, or resource used in the Cloud Services, Software, Products and/or Services, and iii) use, provide, generate, modify, distribute, display and create derivative works of the Outputs; and (d) the AI Services, AI Training Data, and Outputs have been assessed for bias, inaccuracies, unfairness, and controls, mechanisms, and policies are in place to mitigate and ensure the elimination of same. Supplier will reasonably assist and cooperate with Customer to comply with the laws, statutes, and regulations that arise out of or relate to the use of the AI Services by Customer, and regulatory or government entity requests, policies, reporting, or inquiries, that apply to Customer or relate to the Cloud Services, Software, Products and/or Services, AI Services, AI Training Data, or Output.

2.5. **Effect of Termination**. In addition and without limiting Supplier's termination obligations under the Agreement, within thirty (30) calendar days of any termination or expiration of an Agreement or any post-termination transition period (whichever occurs later), Supplier shall (i) provide Customer with all Outputs through the end of the termination effective date, if not already delivered, in a form and format reasonably requested by Customer, and (ii) then destroy the Outputs, and all copies thereof, if not already destroyed, within thirty (30) calendar days after such termination or expiration and provide Customer with written certification of the destruction.

2.6. **Infringement Indemnification**. Without limiting and in addition to Supplier's indemnification obligations under the Agreement, Supplier will defend, indemnify, and hold Customer and KP Affiliates harmless from and against any all

liabilities, damages, claims, costs, expenses (including fees of attorneys and expert witnesses) and other losses resulting from a claim that the AI Training Data, or Outputs furnished to Customer and KP Affiliates, or use thereof, infringe any Intellectual Property Rights of any third party or have become the subject of an injunction or settlement prohibiting the use of such AI Training Data or Outputs, and all infringement remedies in the Agreement will apply to the AI Training Data and Outputs.

2.7. **Additional AI Terms**. Upon Customer's request, Supplier will provide information explaining its governance programs, audits and other mechanisms that ensures the AI technology's usability, reliability, accountability and protections against bias, inaccuracies, and unfairness, including without limitation disclosing any results of bias assessments performed by Supplier, an explanation of how the technology operates, processes data, how the AI Service makes or supports decision making, and any other feature described in an Agreement or that Supplier has described in its documentation, as applicable. Supplier shall provide written notice to Customer of any material changes to the AI model (including model drift) used to provide the AI Services. Additionally, upon Customer's request, Supplier will provide information about how the model or algorithm was trained and developed, how the AI Training Data was collected and used and the type of data therein, the data sources thereto, and whether all necessary rights, permissions, authorizations, and licenses were granted for use of the AI Training Data and Outputs, and data provenance thereof. Supplier shall maintain and comply with all policies and procedures that are compliant with industry standards relating to the ethical, safe, and responsible use, development, and distribution of artificial intelligence technologies. Any rights, title and interests not expressly granted by Customer under this AI Attachment or Agreement are expressly reserved by Customer and KP Affiliates. If Supplier incorporates, uses, or bundles any new or additional AI Services into the Cloud Services, Software, Products and/or Services, or if Supplier provides to Customer or KP Affiliates, or their employees, agents, or contractors use on behalf of Customer or KP Affiliates, any new Supplier AI services, products, or technologies, a) Supplier shall notify Customer in writing, and b) the terms of this AI Attachment will be deemed to apply and shall govern such new or additional AI services, products, or technologies. Such new services, products, or technologies will be deemed to be "AI Services." In the event that any "shrink-wrap," "browse wrap," or "click-wrap" license or other agreement is required to be accepted to access and use the AI Services, the terms of this AI Attachment will nevertheless apply and the terms of the "shrink-wrap," "browse wrap," or "click-wrap" license or other agreement will be null and void. Supplier shall provide Customer with written notice prior to changing any third party artificial intelligence subcontractor that is used to provide the AI Services, and all new subcontractors are subject to the terms of the Agreement and this AI Attachment.

2.8. **M**
provis
writin
natur
oblig

Related Challenges

The table below summarizes key challenges that organizations may face when managing third party relationships.

Challenge	Detail	Possible Mitigation Strategies
Alignment on Versioning	Maintaining consistent alignment with vendors on model version updates is difficult, particularly when updates occur automatically within embedded health IT features.	Require contractual advance notice (e.g., 30 to 90 days) before model version changes. Negotiate version-lock or opt-out rights for clinical applications. Establish internal change management triggers tied to vendor notifications. Note that enterprise vendors typically will not agree to this mitigation.
Unclear Responsibilities	Contracts rarely spell out clinical, operational, and regulatory responsibilities clearly enough for deploying organizations to effectively manage risk.	Develop and require a responsibility assignment matrix (RACI) as a contract resource. Define explicit boundaries for clinical validation, patient safety reporting, and regulatory compliance obligations for each party.
Contractual and Legal Complexity	AI-specific clauses are not yet industry standard, and many legal teams lack sufficient AI expertise to negotiate or review them effectively.	Adopt standard AI terms language. Engage external AI/health IT legal counsel for high-risk procurements. Build internal review checklists with AI-specific red flags.
Vendor Reluctance on Liability	Vendors frequently push back on liability clauses or resist disclosure requirements, citing intellectual property or competitive concerns.	Establish minimum disclosure standards as a condition of vendor consideration, not a negotiating point. Where multiple viable vendors exist, HDOs should consider a structured RFP or "bake-off" to create competitive pressure around contractual terms.
Power Imbalance	Larger organizations typically have more leverage than smaller systems when negotiating terms with AI vendors. Community health centers and critical-access hospitals may accept unfavorable contract language.	For Community Health Centers, join regional or national networks to aggregate negotiating leverage. Use publicly available model contract templates to anchor negotiations.
Burdensome	Audit protocols and customer	Right-size audit requirements using a

Requirements	alignment reports may be too resource-intensive for smaller vendors or for internal teams with limited legal and compliance capacity.	risk-tiered framework (e.g., reserve intensive oversight for high-risk clinical AI). Develop lightweight reporting templates that can be completed by clinical informatics or compliance staff without dedicated legal resources.
Enterprise SaaS & Embedded AI	AI capabilities embedded in enterprise SaaS platforms (e.g., Microsoft 365 Copilot, Google Workspace AI features, EHR-native AI tools) often do not enter through traditional procurement channels and may have no direct vendor relationship to negotiate. This represents a rapidly growing and largely ungoverned category. Organizations may have limited ability to apply standard contract clauses and should rely on platform governance settings, acceptable use policies, and internal controls instead.	Establish a platform AI governance policy that activates regardless of procurement pathway. Inventory embedded AI features across all enterprise platforms. Apply risk categorization to embedded tools to focus oversight where it matters most. Participate in multi-institution coalitions to engage large platform vendors with greater collective leverage.
Post-deployment Non-compliance	Even where pre-contract protections are strong, vendors may fail to meet monitoring, reporting, or remediation obligations after deployment, leaving the organization without a clear path to enforcement.	Adopt a tiered enforcement ladder (notice & cure → remediation plan → access suspension → termination) with defined timelines embedded in every AI vendor contract. Tie enforcement triggers to documented KPI and monitoring breaches.
End User Liability Exposure	When AI informs a decision, residual liability may fall to the staff member making the decision, particularly where AI recommendations are based on incomplete or inaccessible EHR data.	Invest in end user education that emphasizes the limits of AI outputs and the staff obligation to verify AI outputs for the specific context. Coordinate with Subdomain 4.5: Education, Training, and Feedback Playbook.

Next Steps

Links below will connect you to other AI governance playbooks in this series.

Domain 1: AI Policy

Attribution

CHAI AI Governance Playbooks were developed through the contributions and collaboration of more than 150 organizations across the healthcare ecosystem. CHAI extends its sincere appreciation to all participating organizations and contributors for their time, expertise, and thoughtful feedback throughout the development process. We would also like to provide special recognition to the individuals below who requested attribution by name for their contributions to these efforts:

- Josh Hjelmstad, MHA, AVIA Health
- Sandra Gregg, DNP, MHA, RN, PMP, LSSBB, CCMP, HCD, Booz Allen Hamilton, Inc.
- Vidhi Vinay Kothari, Carnegie Mellon University
- Stephen McLaughlin, PhD, Children's National Hospital
- Ryan Marshall Felder, PhD HEC-C, Cleveland Clinic
- Po-Hao Chen, MD, MBA, Cleveland Clinic Foundation
- Katherine Mulready, JD MPS, cliexa
- Holden Groves, MD, Columbia University/NewYork-Presbyterian
- Shakira J. Grant, MBBS, MSCR, CROSS Global Research & Strategy
- Namrata Rastogi, MBBS MRCGP MS, Enigma Ventures
- Aleocidio Sette Balzanelo, MD, MBA, FOLKS
- Rocco Orlando, MD, Hartford HealthCare
- Michael Blumental, HealthLeap
- Romanus M. Joseph, Innovative Health Accountable Care Organization (ACO)
- Rebecca G. Mishuris, MD, MS, MPH, Mass General Brigham
- Nasibeh Zanjirani Farahani, PhD, CSM, Mayo Clinic
- Jason Gillman, MD, FIDSA, MCG Health
- Douglas Graham, Individual Contributor
- Nicole Petroski, Mercy Health Services
- Haris Shuaib, Newton's Tree
- Rémy Ibarcq, Individual Contributor
- Srikar Nallan, MS, Stanford Healthcare
- Jaimie Weber, MD, Tampa General Hospital
- R. Brandon Hunter, MD, FAAP, Texas Children's Hospital
- Angela L. Lipscomb, JD, LL.M, The University of Texas MD Anderson Cancer Center
- Christopher H. Lee, MBA, BSN, RN-BC, UCLA Health
- Abby Steele, Individual Contributor
- Shiba Kuanar, University of Minnesota
- Joshua Miller, JD, CHC, University of Rochester Medicine

- Anne Travis, MD, MSc, Wolters Kluwer
- Michael L. Shaw, JD, ZS

License

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

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

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